

The Multiple Quantile Graphical Model

Almur Ali¹

J. Zico Kolter^{1,2}

Ryan J. Tibshirani^{1,3}

¹Machine Learning Department, Carnegie Mellon University

²Computer Science Department, Carnegie Mellon University

³Department of Statistics, Carnegie Mellon University

Abstract

We introduce the Multiple Quantile Graphical Model (MQGM), which extends the *neighborhood selection* approach of Meinshausen and Bühlmann for learning sparse graphical models. The latter is defined by the basic subproblem of modeling the conditional mean of one variable as a sparse function of all others. Our approach models a set of conditional quantiles of one variable as a sparse function of all others, and hence offers a much richer, more expressive class of conditional distribution estimates. We establish that, under suitable regularity conditions, the MQGM identifies the exact conditional independencies with probability tending to one as the problem size grows, even outside of the usual homoskedastic Gaussian data model. We develop an efficient algorithm for fitting the MQGM using the alternating direction method of multipliers. We also describe a strategy for sampling from the joint distribution that underlies the MQGM estimate. Lastly, we present detailed experiments that demonstrate the flexibility and effectiveness of the MQGM in modeling heteroskedastic non-Gaussian data.

1 Introduction

We consider modeling the joint distribution $\mathbf{Pr}(y_1, \dots, y_d)$ of d random variables, given n independent draws from this distribution $y^{(1)}, \dots, y^{(n)} \in \mathbf{R}^d$, where possibly $d \gg n$. Later, we generalize this setup and consider modeling the conditional distribution $\mathbf{Pr}(y_1, \dots, y_d | x_1, \dots, x_p)$, given n independent pairs $(x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)}) \in \mathbf{R}^{p+d}$. Our starting point is the *neighborhood selection* method of Meinshausen and Bühlmann (2006), which is typically considered in the context of multivariate Gaussian data, and seen as a tool for *covariance selection* (Dempster, 1972): when $\mathbf{Pr}(y_1, \dots, y_d)$ is a multivariate Gaussian distribution, it is a well-known fact y_j and y_k are conditionally independent given the remaining variables if and only if the coefficient corresponding to y_k is zero in the (linear) regression of y_j on all other variables (*e.g.*, Lauritzen (1996)). Therefore, in neighborhood selection we compute, for each $k = 1, \dots, d$, a lasso regression — in order to obtain a small set of conditional dependencies — of y_k on the remaining variables, *i.e.*,

$$\underset{\theta_k \in \mathbf{R}^d}{\text{minimize}} \sum_{i=1}^n \left(y_k^{(i)} - \sum_{j \neq k} \theta_{kj} y_j^{(i)} \right)^2 + \lambda \|\theta_k\|_1, \quad (1)$$

for a tuning parameter $\lambda > 0$. This strategy can be seen as a *pseudolikelihood* approximation (Besag, 1974),

$$\Pr(y_1, \dots, y_d) \approx \prod_{k=1}^d \Pr(y_k | y_{-k}), \quad (2)$$

where y_{-k} denotes all variables except y_k . Under the multivariate Gaussian model for $\Pr(y_1, \dots, y_d)$, the conditional distributions $\Pr(y_k | y_{-k})$, $k = 1, \dots, d$ here are (univariate) Gaussians, and maximizing the pseudolikelihood in (2) is equivalent to separately maximizing the conditionals, as is precisely done in (1) (with induced sparsity), for $k = 1, \dots, d$.

Following the pseudolikelihood-based approach traditionally means carrying out three steps: (i) we write down a suitable family of joint distributions for $\Pr(y_1, \dots, y_d)$, (ii) we derive the conditionals $\Pr(y_k | y_{-k})$, $k = 1, \dots, d$, and then (iii) we maximize each conditional likelihood by (freely) fitting the parameters. Neighborhood selection, and a number of related approaches that came after it (see Section 2.1), can be all thought of in this workflow. In many ways, step (ii) acts as the bottleneck here, and to derive the conditionals, we are usually limited to a homoskedastic and parametric family for the joint distribution.

The approach we take in this paper differs somewhat substantially, as we *begin* by directly modeling the conditionals in (2), without any preconceived model for the joint distribution — in this sense, it may be seen a type of *dependency network* (Heckerman et al., 2000) for continuous data. We also employ heteroskedastic, nonparametric models for the conditional distributions, which allows us great flexibility in learning these conditional relationships. Our method, called the Multiple Quantile Graphical Model (MQGM), is a marriage of ideas in high-dimensional, nonparametric, multiple quantile regression with those in the dependency network literature (the latter is typically focused on discrete, not continuous, data).

An outline for this paper is as follows. Section 2 reviews background material, and Section 3 develops the MQGM estimator. Section 4 studies basic properties of the MQGM, and establishes a structure recovery result under appropriate regularity conditions, even for heteroskedastic, non-Gaussian data. Section 5 describes an efficient ADMM algorithm for estimation, and Section 6 presents empirical examples comparing the MQGM versus common alternatives. Section 7 concludes with a discussion.

2 Background

2.1 Neighborhood selection and related methods

Neighborhood selection has motivated a number of methods for learning sparse graphical models. The literature here is vast; we do not claim to give a complete treatment, but just mention some relevant approaches. Many pseudolikelihood approaches have been proposed, see, *e.g.*, Rocha et al. (2008); Peng et al. (2009); Friedman et al. (2010); Liu and Wang (2012); Khare et al. (2014); Ali et al. (2016). These works exploit the connection between estimating a sparse inverse covariance matrix and regression, and they vary in terms of the optimization algorithms they use and the theoretical guarantees they offer.

In a related but distinct line of research, Yuan and Lin (2007); Banerjee et al. (2008); Friedman et al. (2008); Rothman et al. (2008) proposed ℓ_1 -penalized likelihood estimation in the Gaussian graphical model, a method now generally termed the *graphical lasso* (GLasso).

Following this, several recent papers have extended the GLasso in various ways. Finegold and Drton (2011) examined a modification based on the multivariate Student t -distribution, for robust graphical modeling. Sohn and Kim (2012); Yuan and Zhang (2014); Wytock and Kolter (2013) considered conditional distributions of the form $\mathbf{Pr}(y_1, \dots, y_d | x_1, \dots, x_p)$. Lee and Hastie (2013) proposed a model for mixed (both continuous and discrete) data types, generalizing both GLasso and pairwise Markov random fields. Liu et al. (2009, 2012) used copulas for learning non-Gaussian graphical models.

A strength of neighborhood-based (*i.e.*, pseudolikelihood-based) approaches lies in their simplicity; because they essentially reduce to a collection of univariate probability models, they are in a sense much easier to study outside of the typical homoskedastic, Gaussian data setting. Höfling and Tibshirani (2009); Yang et al. (2012, 2015) elegantly studied the implications of using univariate exponential family models for the conditionals in (2). Closely related to pseudolikelihood approaches are dependency networks (Heckerman et al., 2000). Both frameworks focus on the conditional distributions of one variable given all the rest; the difference lies in whether or not the model for conditionals stems from first specifying some family of joint distributions (pseudolikelihood methods), or not (dependency networks). Dependency networks have been thoroughly studied for discrete data, *e.g.*, Heckerman et al. (2000); Neville and Jensen (2004). For continuous data, Voorman et al. (2014) proposed modeling the mean in a Gaussian neighborhood regression as a nonparametric, additive function of the remaining variables, yielding flexible relationships — this is a type of dependency network for continuous data (though it is not described by the authors in this way). Our method, the MQGM, also deals with continuous data, and is the first to our knowledge that allows for fully nonparametric conditional distributions, as well as nonparametric contributions of the neighborhood variables, in each local model.

2.2 Quantile regression

In linear regression, we estimate the conditional mean of $y | x_1, \dots, x_p$ based on data samples. Similarly, in α -quantile regression (Koenker and Bassett, 1978), we estimate the conditional α -quantile of $y | x_1, \dots, x_p$, formally $Q_{y|x_1, \dots, x_p}(\alpha) = \inf\{t : \mathbf{Pr}(y \leq t | x_1, \dots, x_p) \geq \alpha\}$, for a given $\alpha \in [0, 1]$, by solving the convex optimization problem:

$$\underset{\theta}{\text{minimize}} \quad \sum_{i=1}^n \psi_{\alpha} \left(y^{(i)} - \sum_{j=1}^q \theta_j x_j^{(i)} \right),$$

where $\psi_{\alpha}(z) = \max\{\alpha z, (\alpha - 1)z\}$ is the quantile loss (also referred to as the “pinball” or “tilted absolute” loss). Quantile regression can be useful when the conditional distribution in question is suspected to be heteroskedastic and/or non-Gaussian, *e.g.*, heavy-tailed, or if we wish to understand properties of the distribution other than the mean, *e.g.*, tail behavior. In multiple quantile regression, we solve several quantile regression problems simultaneously, each corresponding to a different quantile level; these problems can be coupled somehow to increase efficiency in estimation (see details in the next section). Again, the literature on quantile regression is quite vast (especially that from econometrics), and we only give a short review here. A standard text is Koenker (2005). Nonparametric modeling of quantiles is a natural extension from the (linear) quantile regression approach outlined above; in the univariate case (one conditioning variable), Koenker et al. (1994) suggested a method

using smoothing splines, and Takeuchi et al. (2006) described an approach using kernels. More recently, Koenker (2011) studied the multivariate nonparametric case (more than one conditioning variable), using additive models. In the high-dimensional setting, where p is large, Belloni and Chernozhukov (2011); Kato (2011); Fan et al. (2014) studied ℓ_1 -penalized quantile regression and derived estimation and recovery theory for non-(sub-)Gaussian data. We extend results in Fan et al. (2014) to prove structure recovery guarantees for the MQGM (in Section 4.3).

3 The multiple quantile graphical model

Many choices can be made with regards to the final form of the MQGM, and to help in understanding these options, we break down our presentation in parts. First fix some ordered set $\mathcal{A} = \{\alpha_1, \dots, \alpha_r\}$ of quantile levels, *e.g.*, $\mathcal{A} = \{0.05, 0.10, \dots, 0.95\}$. For each variable y_k , and each level α_ℓ , we model the α_ℓ -conditional quantile given the other variables, using an additive expansion of the form:

$$Q_{y_k|y_{-k}}(\alpha_\ell) = b_{\ell k}^* + \sum_{j \neq k}^d f_{\ell k j}^*(y_j), \quad (3)$$

where $b_{\ell k}^* \in \mathbf{R}$ is an intercept term, and $f_{\ell k j}^*$, $j = 1, \dots, d$ are smooth, but not parametric in form.

Generic functional form of the MQGM In its most general form, the MQGM estimator is defined as a collection of optimization problems, over $k = 1, \dots, d$ and $\ell = 1, \dots, r$:

$$\underset{b_{\ell k}, f_{\ell k j} \in \mathcal{F}_{\ell k j}, j=1, \dots, d}{\text{minimize}} \sum_{i=1}^n \psi_{\alpha_\ell} \left(y_k^{(i)} - b_{\ell k} - \sum_{j \neq k} f_{\ell k j}(y_j^{(i)}) \right) + \sum_{j \neq k} \left(\lambda_1 P_1(f_{\ell k j}) + \lambda_2 P_2(f_{\ell k j}) \right)^\omega. \quad (4)$$

Here $\lambda_1, \lambda_2 \geq 0$ are tuning parameters. Also, $\mathcal{F}_{\ell k j}$, $j = 1, \dots, d$ are spaces of univariate functions, $\omega > 0$ is a fixed exponent, and P_1, P_2 are sparsity and smoothness penalty functions, respectively, all to be decided as part of the modeling process. We give three examples below; several other variants are possible outside of what we describe.

Example 1: basis expansion model Consider taking $\mathcal{F}_{\ell k j} = \text{span}\{\phi_1^j, \dots, \phi_m^j\}$, the span of m basis functions, *e.g.*, radial basis functions (RBFs) with centers placed at appropriate locations across the domain of variable j , for each $j = 1, \dots, d$. This means that each $f_{\ell k j} \in \mathcal{F}_{\ell k j}$ can be expressed as $f_{\ell k j}(x) = \theta_{\ell k j}^T \phi^j(x)$, for a coefficient vector $\theta_{\ell k j} \in \mathbf{R}^m$, where $\phi^j(x) = (\phi_1^j(x), \dots, \phi_m^j(x))$. Also consider an exponent $\omega = 1$, and the sparsity and smoothness penalties

$$P_1(f_{\ell k j}) = \|\theta_{\ell k j}\|_2 \quad \text{and} \quad P_2(f_{\ell k j}) = \|\theta_{\ell k j}\|_2^2,$$

respectively, which are group lasso and ridge penalties, respectively. With these choices in place, the MQGM problem in (4) can be rewritten in finite-dimensional form:

$$\underset{b_{\ell k}, \theta_{\ell k} = (\theta_{\ell k 1}, \dots, \theta_{\ell k d})}{\text{minimize}} \psi_{\alpha_\ell} \left(Y_k - b_{\ell k} \mathbf{1} - \Phi \theta_{\ell k} \right) + \sum_{j \neq k} \left(\lambda_1 \|\theta_{\ell k j}\|_2 + \lambda_2 \|\theta_{\ell k j}\|_2^2 \right). \quad (5)$$

Above, we used the abbreviation $\psi_{\alpha_\ell}(z) = \sum_{i=1}^n \psi_{\alpha_\ell}(z_i)$ for a vector $z = (z_1, \dots, z_n) \in \mathbf{R}^n$, and also $Y_k = (y_k^{(1)}, \dots, y_k^{(n)}) \in \mathbf{R}^n$ for the observations along variable k , $\mathbf{1} = (1, \dots, 1) \in \mathbf{R}^n$, and $\Phi \in \mathbf{R}^{n \times dm}$ for the basis matrix, with blocks of columns $\Phi_{ij} = \phi(y_j^{(i)})^T \in \mathbf{R}^m$.

The basis expansion model is simple and tends to work well in practice. For the majority of the paper, we will focus on this model; in principle, everything that follows (methodologically, theoretically, algorithmically) extends to the next two models we describe, as well as many other variants.

Example 2: smoothing splines model Now consider taking $\mathcal{F}_{\ell kj} = \text{span}\{g_1^j, \dots, g_n^j\}$, the span of $m = n$ natural cubic splines with knots at $y_j^{(1)}, \dots, y_j^{(n)}$, for $j = 1, \dots, d$. As before, we can then write $f_{\ell kj}(x) = \theta_{\ell kj}^T g^j(x)$ with coefficients $\theta_{\ell kj} \in \mathbf{R}^n$, for $f_{\ell kj} \in \mathcal{F}_{\ell kj}$. The work of Meier et al. (2009), on high-dimensional additive smoothing splines, suggests a choice of exponent $\omega = 1/2$, and penalties

$$P_1(f_{\ell kj}) = \|G^j \theta_{\ell kj}\|_2^2 \quad \text{and} \quad P_2(f_{\ell kj}) = \theta_{\ell kj}^T \Omega^j \theta_{\ell kj},$$

for sparsity and smoothness, respectively, where $G^j \in \mathbf{R}^{n \times n}$ is a spline basis matrix with entries $G_{ii'}^j = g_{i'}^j(y_j^{(i)})$, and Ω^j is the smoothing spline penalty matrix containing integrated products of pairs of twice differentiated basis functions. The MQGM problem in (4) can be translated into a finite-dimensional form, very similar to what we have done in (5), but we omit this for brevity.

Example 3: RKHS model Consider taking $\mathcal{F}_{\ell kj} = \mathcal{H}_j$, a univariate reproducing kernel Hilbert space (RKHS), with kernel function $\kappa^j(\cdot, \cdot)$. The representer theorem allows us to express each function $f_{\ell kj} \in \mathcal{H}_j$ in terms of the representer of evaluation, *i.e.*, $f_{\ell kj}(x) = \sum_{i=1}^n (\theta_{\ell kj})_i \kappa^j(x, y_j^{(i)})$, for a coefficient vector $\theta_{\ell kj} \in \mathbf{R}^n$. The work of Raskutti et al. (2012), on high-dimensional additive RKHS modeling, suggests a choice of exponent $\omega = 1$, and sparsity and smoothness penalties

$$P_1(f_{\ell kj}) = \|K^j \theta_{\ell kj}\|_2 \quad \text{and} \quad P_2(f_{\ell kj}) = \sqrt{\theta_{\ell kj}^T K^j \theta_{\ell kj}},$$

respectively, where $K^j \in \mathbf{R}^{n \times n}$ is the kernel matrix with entries $K_{ii'}^j = \kappa^j(y_j^{(i)}, y_j^{(i')})$. Again, the MQGM problem in (4) can be written in finite-dimensional form, now an SDP, omitted for brevity.

Structural constraints Different kinds of structural constraints can be placed on top of the MQGM optimization problem in order to guide the estimated component functions to meet particular shape requirements. An important example are *non-crossing constraints* (commonplace in nonparametric, multiple quantile regression (Koenker, 2005; Takeuchi et al., 2006)): here, we optimize (4) jointly over $\ell = 1, \dots, r$, subject to

$$b_{\ell k} + \sum_{j \neq k} f_{\ell kj}(y_j^{(i)}) \leq b_{\ell' k} + \sum_{j \neq k} f_{\ell' kj}(y_j^{(i)}), \quad \text{for all } \alpha_\ell < \alpha_{\ell'}, \text{ and } i = 1, \dots, n. \quad (6)$$

This ensures that the estimated quantiles obey the proper ordering, at the observations. For concreteness, we consider the implications for the basis regression model, in Example

1 (similar statements hold for the other two models). For each $\ell = 1, \dots, r$, denote by $F_{\ell k}(b_{\ell k}, \theta_{\ell k})$ the criterion in (5). Introducing the non-crossing constraints requires coupling (5) over $\ell = 1, \dots, r$, so that we now have the following optimization problems, for each target variable $k = 1, \dots, d$:

$$\underset{B_k, \Theta_k}{\text{minimize}} \sum_{\ell=1}^r F_{\ell k}(b_{\ell k}, \theta_{\ell k}) \quad \text{subject to} \quad (\mathbf{1}B_k^T + \Phi\Theta_k)D^T \geq 0, \quad (7)$$

where we denote $B_k = (b_{1k}, \dots, b_{rk}) \in \mathbf{R}^r$, $\Phi \in \mathbf{R}^{n \times dm}$ the basis matrix as before, $\Theta_k \in \mathbf{R}^{dm \times r}$ given by column-stacking $\theta_{\ell k} \in \mathbf{R}^{dm}$, $\ell = 1, \dots, r$, and $D \in \mathbf{R}^{(r-1) \times r}$ is the usual discrete difference operator, *i.e.*,

$$D = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \dots & 0 \\ \vdots & & \ddots & \ddots & \\ 0 & 0 & \dots & -1 & 1 \end{bmatrix}.$$

(The inequality in (7) is to be interpreted componentwise.) Computationally, coupling the subproblems across $\ell = 1, \dots, r$ clearly adds to the overall difficulty of the MQGM, but statistically this coupling acts as a regularizer, by constraining the parameter space in a useful way, thus increasing our efficiency in fitting multiple quantile levels from the given data.

For a triplet ℓ, k, j , *monotonicity constraints* are easy to add, *i.e.*, $f_{\ell k j}(y_j^{(i)}) \leq f_{\ell k j}(y_j^{(i')})$ for all $y_j^{(i)} < y_j^{(i')}$. *Convexity constraints*, where we require $f_{\ell k j}$ to be convex over the observations, for a particular ℓ, k, j , are also straightforward. Lastly, *strong non-crossing constraints*, where we enforce (6) but over all inputs $z \in \mathbf{R}^d$ (not just over the observations) are also possible with positive basis functions.

Exogenous variables and conditional random fields So far, we have considered modeling the joint distribution $\mathbf{Pr}(y_1, \dots, y_d)$, corresponding to learning a Markov random field (MRF). It is not hard to extend our framework to model the conditional distribution $\mathbf{Pr}(y_1, \dots, y_d | x_1, \dots, x_p)$ given some exogenous variables x_1, \dots, x_p , corresponding to learning a conditional random field (CRF). To extend the basis regression model, we introduce the additional parameters $\theta_{\ell k}^x \in \mathbf{R}^p$ in (5), and the loss now becomes $\psi_{\alpha_\ell}(Y_k - b_{\ell k}\mathbf{1}^T - \Phi\theta_{\ell k} - X\theta_{\ell k}^x)$, where $X \in \mathbf{R}^{n \times q}$ is filled with the exogenous observations $x^{(1)}, \dots, x^{(n)} \in \mathbf{R}^q$; the other models are changed similarly.

4 Basic properties and theory

4.1 Quantiles and conditional independence

In the model (3), if a particular variable y_j has no contribution, *i.e.*, satisfied $f_{\ell k j}^* = 0$ across all quantile levels α_ℓ , $\ell = 1, \dots, r$, what does this imply about the conditional independence between y_k and y_j , given the rest? Outside of the multivariate normal model (where the feature transformations need only be linear), nothing can be said in generality. But we argue that conditional independence can be understood in a certain *approximate sense* (*i.e.*, in a

projected approximation of the data generating model). We begin with a simple lemma. Its proof is elementary, and given in the supplement.

Lemma 4.1. *Let U, V, W be random variables, and suppose that all conditional quantiles of $U|V, W$ do not depend on V , i.e., $Q_{U|V, W}(\alpha) = Q_{U|W}(\alpha)$ for all $\alpha \in [0, 1]$. Then U and V are conditionally independent given W .*

By the lemma, if we knew that $Q_{U|V, W}(\alpha) = h(\alpha, U, W)$ for a function h , then it would follow that U, V are conditionally independent given W (*n.b.*, the converse is true, as well). The MQGM problem in (4), with sparsity imposed on the coefficients, essentially aims to achieve such a representation for the conditional quantiles; of course we cannot use a *fully nonparametric* representation of the conditional distribution $y_k|y_{-k}$ and instead we use an *r-step approximation* to the conditional cumulative distribution function (CDF) of $y_k|y_{-k}$ (corresponding to estimating r conditional quantiles), and (say) in the basis regression model, limit the dependence on conditioning variables to be in terms of an additive function of RBFs in y_j , $j \neq k$. Thus, if at the solution in (5) we find that $\hat{\theta}_{kj\ell} = 0$, $\ell = 1, \dots, r$, we may interpret this to mean that y_k and y_j are conditionally independent given the remaining variables, but according to the distribution defined by the *projection* of $y_k|y_{-k}$ onto the space of models considered in (5) (*r*-step conditional CDFs, which are additive expansions in y_j , $j \neq k$). This interpretation is no more tenuous (arguably, less so, as the model space here is much larger) than that needed when applying standard neighborhood selection to non-Gaussian data.

4.2 Gibbs sampling and the “joint” distribution

When specifying a form for the conditional distributions in a pseudolikelihood approximation as in (2), it is natural to ask: what is the corresponding joint distribution? Unfortunately, for a general collection of conditional distributions, there need not exist a compatible joint distribution, even when all conditionals are continuous (Wang and Ip, 2008). Still, pseudolikelihood approximations (a special case of composite likelihood approximations), possess solid theoretical backing, in that maximizing the pseudolikelihood relates closely to minimizing a certain (expected composite) Kullback-Leibler divergence, measured to the true conditionals (Varin and Vidoni, 2005). Recently, Chen et al. (2015); Yang et al. (2015) made nice progress in describing specific conditions on conditional distributions that give rise to a valid joint distribution, though their work was specific to exponential families. A practical answer to the question of this subsection is to use Gibbs sampling, which attempts to draw samples consistent with the fitted conditionals; this is precisely the observation of Heckerman et al. (2000), who show that Gibbs sampling from discrete conditionals converges to a unique stationary distribution, although this distribution may not actually be compatible with the conditionals. The following result establishes the analogous claim for continuous conditionals; its proof is in the supplement. We demonstrate the practical value of Gibbs sampling through various examples in Section 6.

Lemma 4.2. *Assume that the conditional distributions $\Pr(y_k|y_{-k})$, $k = 1, \dots, d$ take only positive values on their domain. Then, for any given ordering of the variables, Gibbs sampling converges to a unique stationary distribution that can be reached from any initial point. (This stationary distribution depends on the ordering.)*

4.3 Graph structure recovery

When $\log d = O(n^{2/21})$, and we assume somewhat standard regularity conditions (listed as A1–A4 in the supplement), we will show that the MQGM estimate recovers the underlying conditional independencies with high probability (interpreted in the projected model space, as explained in Section 4.1). Importantly, we do not require a Gaussian, sub-Gaussian, or even parametric assumption on the data generating process; instead, we assume i.i.d. draws $y^{(1)}, \dots, y^{(n)} \in \mathbf{R}^d$, where the conditional distributions $y_k | y_{-k}$ have quantiles that are specified by the model in (3) for $k = 1, \dots, d$, $\ell = 1, \dots, r$, and further, each $f_{\ell kj}^*(x) = \theta_{\ell kj}^T \phi^j(x)^*$ for coefficients $\theta_{\ell kj}^* \in \mathbf{R}^m$, $j = 1, \dots, d$, as in the basis expansion model.

Let E^* denote the corresponding edge set of conditional dependencies from these neighborhood models, *i.e.*, $\{k, j\} \in E^* \iff \max_{\ell=1, \dots, r} \max\{\|\theta_{\ell kj}^*\|_2, \|\theta_{\ell jk}^*\|_2\} > 0$. We define the estimated edge set \hat{E} in the analogous way, based on the solution in (5). Without a loss of generality, we assume the features have been scaled to satisfy $\|\Phi_j\| \leq \sqrt{n}$ for $j = 1, \dots, dm$. The following is our recovery result; its proof is provided in the supplement.

Theorem 4.3. *Assume $\log d = O(n^{2/21})$, and conditions A1–A4 in the supplement. Assume that the tuning parameters λ_1, λ_2 satisfy*

$$\lambda_1 \asymp \sqrt{mn \log(d^2 mr / \delta) \log^3 n} \quad \text{and} \quad \lambda_2 = o(n^{41/42} / \theta_{\max}^*),$$

where $\theta_{\max}^* = \max_{\ell, k, j} \|\theta_{\ell kj}^*\|_2$. Then for n large enough, the MQGM estimate in (5) exactly recovers the underlying conditional dependencies, *i.e.*, $\hat{E} = E^*$, with probability at least $1 - \delta$.

The theorem shows that the nonzero pattern in the MQGM estimate identifies, with high probability, the underlying conditional independencies. But to be clear, we emphasize that the MQGM estimate is *not* an estimate of the inverse covariance matrix itself (this is also the case with neighborhood regression, SpaceJam of Voorman et al. (2014), and many other methods for learning graphical models).

5 Computational approach

By design, the MQGM problem in (5) separates into d subproblems, across $k = 1, \dots, d$ (it therefore suffices to consider only a single subproblem, so we omit notational dependence on k for auxiliary variables). While these subproblems are challenging for off-the-shelf solvers (even for only moderately-sized graphs), the key terms here all admit efficient *proximal operators* (Parikh and Boyd, 2013), which makes operator splitting methods like the alternating direction method of multipliers (Boyd et al., 2011) a natural choice. As an illustration, we consider the non-crossing constraints in the basis regression model below. Reparameterizing so that we may apply ADMM:

$$\begin{aligned} & \underset{\Theta_k, B_k, V, W, Z}{\text{minimize}} && \psi_{\mathcal{A}}(Z) + \lambda_1 \sum_{\ell=1}^r \sum_{j=1}^d \|W_{\ell j}\|_2 + \frac{\lambda_2}{2} \|W\|_F^2 + I_+(VD^T) \\ & \text{subject to} && V = \mathbf{1} B_k^T + \Phi \Theta_k, \quad W = \Theta_k, \quad Z = Y_k \mathbf{1}^T - \mathbf{1} B_k^T - \Phi \Theta_k, \end{aligned} \tag{8}$$

where for brevity $\psi_{\mathcal{A}}(A) = \sum_{\ell=1}^r \sum_{j=1}^d \psi_{\alpha_{\ell}}(A_{\ell j})$, and $I_+(\cdot)$ is the indicator function of the space of elementwise nonnegative matrices. The augmented Lagrangian associated with (8)

is:

$$L_\rho(\Theta_k, B_k, V, W, Z, U_V, U_W, U_Z) = \psi_{\mathcal{A}}(Z) + \lambda_1 \sum_{\ell=1}^r \sum_{j=1}^d \|W_{\ell j}\|_2 + \frac{\lambda_2}{2} \|W\|_F^2 + I_+(VD^T) \\ + \frac{\rho}{2} \left(\|\mathbf{1}B_k^T + \Phi\Theta_k - V + U_V\|_F^2 + \|\Theta_k - W + U_W\|_F^2 + \|Y_k\mathbf{1}^T - \mathbf{1}B_k^T - \Phi\Theta_k - Z + U_Z\|_F^2 \right), \quad (9)$$

where $\rho > 0$ is the augmented Lagrangian parameter, and U_V, U_W, U_Z are dual variables corresponding to the equality constraints on V, W, Z , respectively. Minimizing (9) over V yields:

$$V \leftarrow P_{\text{iso}}(\mathbf{1}B_k^T + \Phi\Theta_k + U_V), \quad (10)$$

where $P_{\text{iso}}(\cdot)$ denotes the row-wise projection operator onto the isotonic cone (the space of componentwise nondecreasing vectors), an $O(nr)$ operation here (Johnson, 2013). Minimizing (9) over $W_{\ell j}$ yields the update:

$$W_{\ell j} \leftarrow \frac{(\Theta_k)_{\ell j} + (U_W)_{\ell j}}{1 + \lambda_2/\rho} \left(1 - \frac{\lambda_1/\rho}{\|(\Theta_k)_{\ell j} + (U_W)_{\ell j}\|_2} \right)_+, \quad (11)$$

where $(\cdot)_+$ is the positive part operator. This can be seen by deriving the proximal operator of the function $f(x) = \lambda_1\|x\|_2 + (\lambda_2/2)\|x\|_2^2$. Minimizing (9) over Z yields the update:

$$Z \leftarrow \mathbf{prox}_{(1/\rho)\psi_{\mathcal{A}}}(Y_k\mathbf{1}^T - \mathbf{1}b_k^T - \Phi\Theta_k + U_Z), \quad (12)$$

where $\mathbf{prox}_f(\cdot)$ denotes the proximal operator of a function f . For the multiple quantile loss function $\psi_{\mathcal{A}}$, this is a kind of generalized soft-thresholding. The proof is given in the supplement.

Lemma 5.1. *Let $P_+(\cdot)$ and $P_-(\cdot)$ be the elementwise positive and negative part operators, respectively, and let $a = (\alpha_1, \dots, \alpha_r)$. Then $\mathbf{prox}_{t\psi_{\mathcal{A}}}(A) = P_+(A - t\mathbf{1}a^T) + P_-(A - t\mathbf{1}a^T)$.*

Finally, differentiation in (9) with respect to B_k and Θ_k yields the simultaneous updates:

$$\begin{bmatrix} \Theta_k \\ B_k^T \end{bmatrix} \leftarrow \frac{1}{2} \begin{bmatrix} \Phi^T\Phi + \frac{1}{2}I & \Phi^T\mathbf{1} \\ \mathbf{1}^T\Phi & \mathbf{1}^T\mathbf{1} \end{bmatrix}^{-1} \left([I \ \mathbf{0}]^T (W - U_W) + [\Phi \ \mathbf{1}]^T (Y_k\mathbf{1}^T - Z + U_Z + V - U_V) \right). \quad (13)$$

A complete description of our ADMM algorithm for solving the MQGM problem is in the supplement.

Gibbs sampling Having fit the conditionals $y_k|y_{-k}$, $k = 1, \dots, d$, we may want to make predictions or extract joint distributions over subsets of variables. As discussed in Section 4.2, there is no general analytic form for these joint distributions, but the pseudolikelihood approximation underlying the MQGM suggests a natural Gibbs sampler. A careful implementation that respects the additive model in (3) yields a highly efficient Gibbs sampler, especially for CRFs; the supplement gives details.

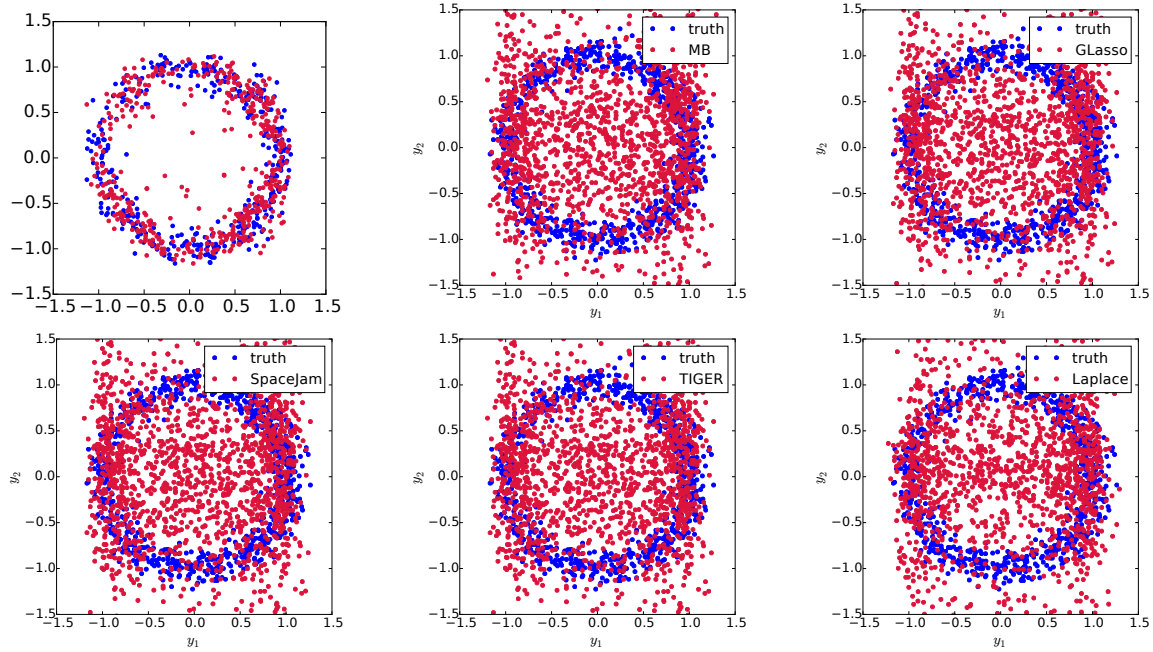


Figure 1: Data from the ring distribution (blue) as well as new samples (red) from the MQGM, MB, GLasso, SpaceJam, TIGER, and Laplace fitted to the same (blue) data; the samples from the MQGM were obtained by using our Gibbs sampler.

6 Empirical examples

6.1 Synthetic data

We consider synthetic examples, comparing to neighborhood selection (MB), the graphical lasso (GLasso), SpaceJam Voorman et al. (2014), the nonparanormal skeptic Liu et al. (2012), TIGER Liu and Wang (2012), and neighborhood selection using the absolute loss (Laplace).

Ring example As a simple but telling example, we drew $n = 400$ samples from a “ring” distribution in $d = 4$ dimensions. We used $m = 10$ expanded features and $r = 20$ quantile levels. Data were generated by first drawing a random angle $\nu \sim \text{Uniform}(0, 1)$, then a random radius $R \sim \mathcal{N}(0, 0.1)$, and finally computing the coordinates $y_1 = R \cos \nu$, $y_2 = R \sin \nu$ and $y_3, y_4 \sim \mathcal{N}(0, 1)$, *i.e.*, y_1 and y_2 are the only dependent variables here. Figure 1 plots samples (blue) of the coordinates (y_1, y_2) as well as new samples (red) from the MQGM, MB, GLasso, and SpaceJam fitted to these same (blue) samples; the samples from the MQGM, obtained by using our Gibbs sampler (see the supplement for further details), appear to closely match the samples from the underlying ring.

Figure 2 shows the conditional independencies recovered by the MQGM, MB, GLasso, SpaceJam, the nonparanormal skeptic, TIGER, and Laplace, when run on the ring data. We visualize these independencies by forming a $d \times d$ matrix with the cell (j, k) set to white if j, k are conditionally independent given the others, and black otherwise. Across a range of tuning parameters for each method, the MQGM is the only one that successfully recovers




Figure 2: Conditional independencies recovered by the MQGM, MB, GLasso, SpaceJam, the nonparanormal skeptic, TIGER, and Laplace on the ring data; black means conditional dependence. The MQGM is the only method that successfully recovers the underlying conditional dependencies.

the underlying conditional dependencies.

Table 1 presents an evaluation of the conditional CDFs given by the MQGM, MB, GLasso, SpaceJam, TIGER, and Laplace when run on the ring data. For each method, we averaged the total variation distances and Kolmogorov-Smirnoff statistics between the fitted and true conditional CDFs across all variables, and then reported the best values obtained across a range of tuning parameters (further details are given in the supplement); the MQGM outperforms all its competitors, in both metrics.

Larger examples To investigate performance at larger scales, we drew $n \in \{50, 100, 300\}$ samples from a multivariate normal and Student t -distribution (with 3 degrees of freedom), both in $d = 100$ dimensions, and parameterized by a random, sparse, diagonally dominant $d \times d$ inverse covariance matrix, following the procedure in Peng et al. (2009); Khare et al. (2014); Oh et al. (2014); Ali et al. (2016). Over the same set of sample sizes, with $d = 100$, we also considered an autoregressive setup in which we drew samples of pairs of adjacent variables from the ring distribution. In all three data settings (normal, t , and autoregressive), we used $m = 10$ and $r = 20$ for the MQGM. To summarize the performances, we considered a

Table 1: Total variation (TV) distance and Kolmogorov-Smirnoff (KS) statistic values for the MQGM, MB, GLasso, SpaceJam, TIGER, and Laplace on the ring data; lower is better, best in **bold**.

	TV	KS
MQGM	20.873	0.760
MB	92.298	1.856
GLasso	92.479	1.768
SpaceJam	91.568	1.697
TIGER	88.284	1.450
Laplace	127.406	1.768

Table 2: AUC values for the MQGM, MB, GLasso, SpaceJam, the nonparanormal skeptic, TIGER, and Laplace for the normal, t , and autoregressive data settings; higher is better, best in **bold** (standard errors are $\approx 10^{-4}$ or smaller).

	Normal			Student t			Autoregressive		
	$n = 50$	$n = 100$	$n = 300$	$n = 50$	$n = 100$	$n = 300$	$n = 50$	$n = 100$	$n = 300$
MQGM	0.953	0.976	0.988	0.928	0.947	0.981	0.726	0.754	0.955
MB	0.850	0.959	0.994	0.844	0.923	0.988	0.532	0.563	0.725
GLasso	0.908	0.964	0.998	0.691	0.605	0.965	0.541	0.620	0.711
SpaceJam	0.889	0.968	0.997	0.893	0.965	0.993	0.624	0.708	0.854
Nonpara.	0.881	0.962	0.996	0.862	0.942	0.998	0.545	0.590	0.612
TIGER	0.732	0.921	0.996	0.420	0.873	0.989	0.503	0.518	0.718
Laplace	0.803	0.931	0.989	0.800	0.876	0.991	0.530	0.554	0.758

range of tuning parameters for each method, computed corresponding false and true positive rates (in detecting conditional dependencies), and then computed the corresponding area under the curve (AUC), following, *e.g.*, Peng et al. (2009); Khare et al. (2014); Oh et al. (2014); Ali et al. (2016). Table 2 reports the median AUCs (across 50 trials) for all three of these examples; the MQGM outperforms all other methods on the autoregressive example, as well as on the small- n normal and Student t examples.

6.2 Modeling flu epidemics

We study $n = 937$ weekly flu incidence reports from September 28, 1997 through August 30, 2015, across 10 regions in the United States (see the left panel of Figure 3), obtained from Centers for Disease Control and Prevention (CDC) (2015). We considered $d = 20$ variables: the first 10 encode the current week’s flu incidence (precisely, the percentage of doctor’s visits in which flu-like symptoms are presented) in the 10 regions, and the last 10 encode the same but for the prior week. We set $m = 5$, $r = 99$, and also introduced exogenous variables to encode the week numbers, so $p = 1$. Thus, learning the MQGM here corresponds to learning the structure of a spatiotemporal graphical model, and reduces to solving 20 multiple quantile regression subproblems, each of dimension $(19 \times 5 + 1) \times 99 = 9504$. All subproblems took about 1 minute on a 6 core 3.3 Ghz Core i7 X980 processor.

The left panel of Figure 4 plots the wallclock time (seconds) for solving one subproblem using ADMM versus SCS (O’Donoghue et al., 2013), a cone solver that has been advocated as a reasonable choice for a class of problems encapsulating (4); ADMM outperforms SCS by roughly two orders of magnitude. The right panel of Figure 3 presents the conditional

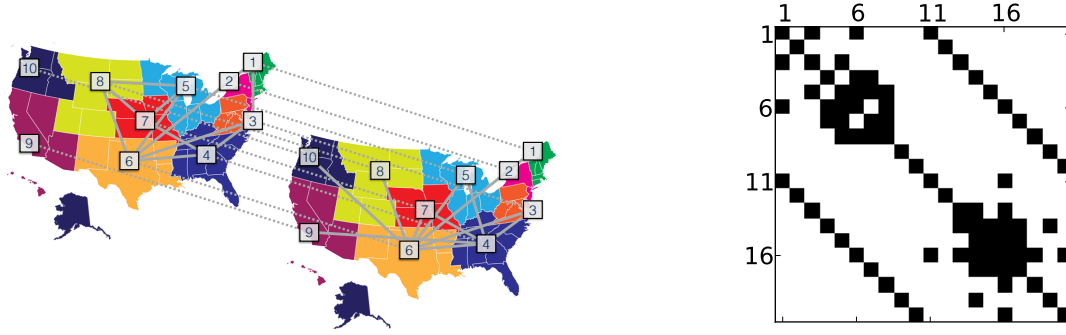


Figure 3: Conditional dependencies recovered by the MQGM on the flu data; each of the first ten cells on the right corresponds to a region of the U.S. on the left, and black means dependence.




Figure 4: Wallclock time (seconds) for solving one subproblem using ADMM versus SCS, on the left. Samples from the fitted marginal distribution of the weekly flu incidence rates at region six, on the right. Samples at larger quantiles are shaded lighter; the median is in darker blue.

independencies recovered by the MQGM. Nonzero entries in the upper left 10×10 submatrix correspond to dependencies between the y_k variables for $k = 1, \dots, 10$; *e.g.*, the nonzero (0,2) entry suggests that region 1 and 3's flu reports are dependent. The lower right 10×10 submatrix corresponds to the y_k variables for $k = 11, \dots, 20$, and the nonzero banded entries suggest that at any region the previous week's flu incidence (naturally) influences the next week's. The left panel of Figure 3 visualizes these relationships by drawing an edge between dependent regions; region 6 is highly connected, suggesting that it is a bellwether for other regions, which is a qualitative observation also made by the CDC. To draw samples from the fitted distributions, we ran our Gibbs sampler over the year, generating 1000 total samples, making 5 passes over all coordinates between each sample, and with a burn-in period of 100 iterations. The right panel of Figure 4 plots samples from the marginal distribution of the percentages of flu reports at region six (other regions are in the supplement) throughout the year, revealing the heteroskedastic nature of the data; we also see that flu incidence (naturally) increases towards the end of the year.

Our last example, on wind power data, is presented in the supplement.

7 Discussion

We proposed and studied the Multiple Quantile Graphical Model (MQGM). We established theoretical and empirical backing to the claim that the MQGM is capable of compactly representing relationships between heteroskedastic, non-Gaussian variables. We developed efficient algorithms for estimation and sampling in the MQGM. All in all, we believe that our work represents a step forward in the design of flexible yet tractable graphical models.

Acknowledgments

This work was supported by a Department of Energy Computational Science Graduate Fellowship, under grant number DE-FG02-97ER25308.

References

- Alnur Ali, Kshitij Khare, Sang-Yun Oh, and Bala Rajaratnam. Generalized pseudolikelihood methods for inverse covariance estimation. Technical report, 2016. Available at <http://arxiv.org/pdf/1606.00033.pdf>.
- Onureena Banerjee, Laurent El Ghaoui, and Alexandre d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516, 2008.
- Alexandre Belloni and Victor Chernozhukov. ℓ_1 -penalized quantile regression in high-dimensional sparse models. *Annals of Statistics*, 39(1):82–130, 2011.
- Julian Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B*, 36(2):192–236, 1974.
- Christopher Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- Centers for Disease Control and Prevention (CDC). Influenza national and regional level graphs and data, August 2015. URL <http://gis.cdc.gov/grasp/fluview/fluportaldashboard.html>.
- Shizhe Chen, Daniela Witten, and Ali Shojaie. Selection and estimation for mixed graphical models. *Biometrika*, 102(1):47–64, 2015.
- Arthur Dempster. Covariance selection. *Biometrics*, 28(1):157–175, 1972.
- Jianqing Fan, Yingying Fan, and Emre Barut. Adaptive robust variable selection. *Annals of Statistics*, 42(1):324–351, 2014.
- Michael Finegold and Mathias Drton. Robust graphical modeling of gene networks using classical and alternative t-distributions. *Annals of Applied Statistics*, 5(2A):1057–1080, 2011.

- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Applications of the lasso and grouped lasso to the estimation of sparse graphical models. Technical report, 2010. Available at <http://statweb.stanford.edu/~tibs/ftp/ggraph.pdf>.
- Laszlo Gyrofi, Michael Kohler, Adam Krzyzak, and Harro Walk. *A Distribution-Free Theory of Nonparametric Regression*. 2002.
- David Heckerman, David Maxwell Chickering, David Meek, Robert Rounthwaite, and Carl Kadie. Dependency networks for inference, collaborative filtering, and data visualization. *Journal of Machine Learning Research*, 1:49–75, 2000.
- Holger Höfling and Robert Tibshirani. Estimation of sparse binary pairwise Markov networks using pseudo-likelihoods. *Journal of Machine Learning Research*, 10:883–906, 2009.
- Tao Hong, Pierre Pinson, and Shu Fan. Global energy forecasting competition 2012. *International Journal of Forecasting*, 30:357–363, 2014.
- Nicholas Johnson. A dynamic programming algorithm for the fused lasso and ℓ_0 -segmentation. *Journal of Computational and Graphical Statistics*, 22(2):246–260, 2013.
- Kengo Kato. Group lasso for high dimensional sparse quantile regression models. Technical report, 2011. Available at <http://arxiv.org/pdf/1103.1458.pdf>.
- Kshitij Khare, Sang-Yun Oh, and Bala Rajaratnam. A convex pseudolikelihood framework for high dimensional partial correlation estimation with convergence guarantees. *Journal of the Royal Statistical Society: Series B*, 77(4):803–825, 2014.
- Roger Koenker. *Quantile Regression*. Cambridge University Press, 2005.
- Roger Koenker. Additive models for quantile regression: Model selection and confidence band-aids. *Brazilian Journal of Probability and Statistics*, 25(3):239–262, 2011.
- Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978.
- Roger Koenker, Pin Ng, and Stephen Portnoy. Quantile smoothing splines. *Biometrika*, 81(4):673–680, 1994.
- Steffen Lauritzen. *Graphical models*. Oxford University Press, 1996.
- Jason Lee and Trevor Hastie. Structure learning of mixed graphical models. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, pages 388–396, 2013.
- Han Liu and Lie Wang. TIGER: A tuning-insensitive approach for optimally estimating Gaussian graphical models. Technical report, 2012. Available at <http://arxiv.org/pdf/1209.2437.pdf>.

- Han Liu, John Lafferty, and Larry Wasserman. The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *Journal of Machine Learning Research*, 10:2295–2328, 2009.
- Han Liu, Fang Han, Ming Yuan, John Lafferty, and Larry Wasserman. High-dimensional semiparametric Gaussian copula graphical models. *The Annals of Statistics*, pages 2293–2326, 2012.
- Lukas Meier, Sara van de Geer, and Peter Bühlmann. High-dimensional additive modeling. *Annals of Statistics*, 37(6):3779–3821, 2009.
- Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics*, 34(3):1436–1462, 2006.
- Jennifer Neville and David Jensen. Dependency networks for relational data. In *Proceedings of Fourth IEEE International Conference on the Data Mining*, pages 170–177. IEEE, 2004.
- Brendan O’Donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. Operator splitting for conic optimization via homogeneous self-dual embedding. Technical report, 2013. Available at <https://stanford.edu/~boyd/papers/pdf/scs.pdf>.
- Sang-Yun Oh, Onkar Dalal, Kshitij Khare, and Bala Rajaratnam. Optimization methods for sparse pseudolikelihood graphical model selection. In *Advances in Neural Information Processing Systems 27*, pages 667–675, 2014.
- Neal Parikh and Stephen Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.
- Jie Peng, Pei Wang, Nengfeng Zhou, and Ji Zhu. Partial correlation estimation by joint sparse regression models. *Journal of the American Statistical Association*, 104(486):735–746, 2009.
- Garvesh Raskutti, Martin Wainwright, and Bin Yu. Minimax-optimal rates for sparse additive models over kernel classes via convex programming. *Journal of Machine Learning Research*, 13:389–427, 2012.
- Guilherme Rocha, Peng Zhao, and Bin Yu. A path following algorithm for sparse pseudolikelihood inverse covariance estimation (SPLICE). Technical report, 2008. Available at <https://www.stat.berkeley.edu/~binyu/ps/rocha.pseudo.pdf>.
- Adam Rothman, Peter Bickel, Elizaveta Levina, and Ji Zhu. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.
- Kyung-Ah Sohn and Seyoung Kim. Joint estimation of structured sparsity and output structure in multiple-output regression via inverse covariance regularization. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics*, pages 1081–1089, 2012.
- Ichiro Takeuchi, Quoc Le, Timothy Sears, and Alexander Smola. Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7:1231–1264, 2006.

- Cristiano Varin and Paolo Vidoni. A note on composite likelihood inference and model selection. *Biometrika*, 92(3):519–528, 2005.
- Arend Voorman, Ali Shojaie, and Daniela Witten. Graph estimation with joint additive models. *Biometrika*, 101(1):85–101, 2014.
- Martin Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (Lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, 2009.
- Yuchung Wang and Edward Ip. Conditionally specified continuous distributions. *Biometrika*, 95(3):735–746, 2008.
- Matt Wytock and Zico Kolter. Sparse Gaussian conditional random fields: Algorithms, theory, and application to energy forecasting. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1265–1273, 2013.
- Eunho Yang, Pradeep Ravikumar, Genevera Allen, and Zhandong Liu. Graphical models via generalized linear models. In *Advances in Neural Information Processing Systems 25*, pages 1358–1366, 2012.
- Eunho Yang, Pradeep Ravikumar, Genevera Allen, and Zhandong Liu. Graphical models via univariate exponential family distributions. *Journal of Machine Learning Research*, 16: 3813–3847, 2015.
- Ming Yuan and Yi Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.
- Xiao-Tong Yuan and Tong Zhang. Partial Gaussian graphical model estimation. *IEEE Transactions on Information Theory*, 60(3):1673–1687, 2014.

Supplement to “The Multiple Quantile Graphical Model”

Almur Ali¹

J. Zico Kolter^{1,2}

Ryan J. Tibshirani^{1,3}

¹Machine Learning Department, Carnegie Mellon University

²Computer Science Department, Carnegie Mellon University

³Department of Statistics, Carnegie Mellon University

S.1 Proof of Lemma 4.1

If the conditional quantiles satisfy $Q_{U|V,W}(\alpha) = Q_{U|W}(\alpha)$ for all $\alpha \in [0, 1]$, then the conditional CDF must obey the same property, *i.e.*, $F_{U|V,W}(t) = F_{U|W}(t)$ for all t in the support of U . This is simply because any CDF may be expressed in terms of its corresponding quantile function (*i.e.*, inverse CDF), as in

$$F_{U|V,W}(t) = \sup\{\alpha \in [0, 1] : Q_{U|V,W}(\alpha) \leq t\},$$

and the right-hand side does not depend on V , so neither can the left-hand side. But this precisely implies that the distribution of $U|V, W$ equals that of $U|W$, *i.e.*, U and V are conditionally independent given W . We note that the converse of the statement in the lemma is true as well, by just reversing all the arguments here. \square

S.2 Proof of Lemma 4.2

This result can be seen as a generalization of Theorem 3 in Heckerman et al. (2000).

First, we define an *iteration* of Gibbs sampling to be a single pass through all the variables (without a loss of generality, we take this order to be y_1, \dots, y_d). Now, consider a particular iteration of Gibbs sampling; let $\tilde{y}_1, \dots, \tilde{y}_d$ be the values assigned to the variables on the previous iteration. Then the transition kernel for our Gibbs sampler is given by

$$\mathbf{Pr}(y_1, \dots, y_d | \tilde{y}_1, \dots, \tilde{y}_d) = \mathbf{Pr}(y_d | y_{d-1}, \dots, y_1, \tilde{y}_1, \dots, \tilde{y}_d) \mathbf{Pr}(y_{d-1}, \dots, y_1 | \tilde{y}_1, \dots, \tilde{y}_d) \quad (\text{S.1})$$

$$= \mathbf{Pr}(y_d | y_{d-1}, \dots, y_1) \mathbf{Pr}(y_{d-1}, \dots, y_1 | \tilde{y}_1, \dots, \tilde{y}_d) \quad (\text{S.2})$$

$$= \mathbf{Pr}(y_d | y_{d-1}, \dots, y_1) \mathbf{Pr}(y_{d-1} | y_{d-2}, \dots, y_1, \tilde{y}_d) \cdots \mathbf{Pr}(y_1 | \tilde{y}_2, \dots, \tilde{y}_d), \quad (\text{S.3})$$

where (S.1) follows by the definition of conditional probability, (S.2) by conditional independence, and (S.3) by repeated applications of these tools. Since each conditional distribution is assumed to be (strictly) positive, we have that the transition kernel is also positive, which in turn implies (Bishop, 2006, page 544) that the induced Markov chain is ergodic with a unique stationary distribution that can be reached from any initial point. \square

S.3 Statement and discussion of regularity conditions for Theorem 4.3

For each $k = 1, \dots, r$, $\ell = 1, \dots, r$, let us define the “effective” (independent) error terms $\epsilon_{\ell ki} = y_k^{(i)} - b_{\ell k}^* - \sum_{j \neq k} \phi(y_j^{(i)})^T \theta_{\ell kj}^*$, over $i = 1, \dots, n$. Denote by $F_{\epsilon_{\ell k}}$ the conditional CDF of $\epsilon_{\ell ki} | y_{-k}^{(i)}$, $i = 1, \dots, n$, which by construction satisfies $F_{\epsilon_{\ell k}}(0) = \alpha_\ell$. Also define the underlying support

$$S_{\ell k} = \{j \in \{1, \dots, d\} : \theta_{\ell kj}^* \neq 0\}.$$

Here we take a moment to explain a somewhat subtle indexing issue with the columns of the feature matrix $\Phi \in \mathbf{R}^{n \times dm}$. For a single fixed index $j = 1, \dots, d$, we will extract an appropriate block of columns of $\Phi \in \mathbf{R}^{n \times dm}$, corresponding to the basis expansion of variable j , by writing Φ_j . More precisely, we use Φ_j to denote the block of m columns

$$[\Phi_{(j-1)m+1}, \Phi_{(j-1)m+2}, \dots, \Phi_{jm}]. \quad (\text{S.4})$$

We do this because it simplifies notation considerably. (Occasionally, to be transparent, we will use the more exhaustive notation on the right-hand side in (S.4), but this is to be treated as an exception, and the default is to use the concise notation as in Φ_j .) The same rule will be used for subsets of indices among $1, \dots, d$, so that $\Phi_{S_{\ell k}}$ denotes the appropriate block of $m|S_{\ell k}|$ columns corresponding to the basis expansions of the variables in $S_{\ell k}$.

For all $k = 1, \dots, d$, $\ell = 1, \dots, r$, we will assume the following regularity conditions.

- A1.** *Groupwise irrerepresentability:* for $j \in S_{\ell k}^c$, we require that $\|\Phi_j^T \Phi_{S_{\ell k}}\|_F < \lambda_1 / (6f_{\epsilon_{\ell k}}(0)\gamma)$, where $S_{\ell k} = \{j \in \{1, \dots, dm\} : \theta_{\ell kj}^* \neq 0\}$, $f_{\epsilon_{\ell k}}$ is the density of $F_{\epsilon_{\ell k}}$, and $\gamma > 0$ is a quantity prescribed by Lemma S.5.
- A2.** *Distributional smoothness:* we assume that $|F_{\epsilon_{\ell k}}(x) - F_{\epsilon_{\ell k}}(0) - xf_{\epsilon_{\ell k}}(0)| \leq C_1 x^2$ for all $|x| \leq C_2$, where $C_1, C_2 > 0$ are constants.
- A3.** *Correlation restriction:* we assume that $C_3 \leq (f_{\epsilon_{\ell k}}(0)/n)\lambda_{\min}(\Phi_{S_{\ell k}}^T \Phi_{S_{\ell k}}) \leq C_4$ for constants $C_3, C_4 > 0$, where $\lambda_{\min}(A)$ denotes the minimum eigenvalue of A .
- A4.** *Basis and support size restrictions:* we assume that $m = O(n^{1/9})$ and $s = O(n^{1/21})$, where $s = |S_{\ell k}|$. We also assume, with probability tending to one, that $\Phi_{\max} = \Omega(1)$ and $\Phi_{\max} = o(n^{1/21} / \log^{1/2} n)$, where we write Φ_{\max} to denote the maximum absolute entry of the basis matrix Φ .

Next, we provide some intuition for these conditions.

Condition A1. Fix some $j \in S_{\ell k}^c$. For notational convenience, we let

$$A = \Phi_j^T \Phi_{S_{\ell k}} \in \mathbf{R}^{m \times sm}.$$

Observe that each entry of A can be expressed as

$$A_{ip} = n\rho_{i,p} \|\Phi_{(j-1)m+i}\|_2 \|\Phi_p\|_2, \quad (\text{S.5})$$

for $i = 1, \dots, m$, p denoting an index into the basis expansion of the columns $\Phi_{S_{\ell k}}$, and $\rho_{i,p}$ denoting the sample correlation coefficient for the columns Φ_i and Φ_p . Since $\|A_p\|_F \leq \sqrt{m}\|A_p\|_\infty$, we have that

$$\max_{i,p} \rho_{i,p} < \frac{\lambda_1}{6n^2 f_{\epsilon_{\ell k}}(0) \sqrt{m}}$$

is sufficient for condition A1; here, we have also used the column scaling assumption $\|\Phi_p\|_2 \leq \sqrt{n}$.

So, roughly speaking, bounded correlation between each pair of columns in the submatrices Φ_j and $\Phi_{S_{\ell k}}$ is enough for condition A1 to hold; note that this is trivially satisfied when $\Phi_i^T \Phi_p = 0$, for $i = 1, \dots, m$, and p as defined above. Condition A1 is therefore similar to, *e.g.*, the mutual incoherence condition of Wainwright (2009) for the lasso, which is given by

$$\left\| \Phi_{S^c}^T \Phi_S (\Phi_S^T \Phi_S)^{-1} \right\|_\infty \leq 1 - \tilde{\gamma} \iff \max_{j \in S^c} \left\| (\Phi_S^T \Phi_S)^{-1} \Phi_S^T \Phi_j \right\|_1 \leq 1 - \tilde{\gamma},$$

where again Φ_S extracts the appropriate block of columns of Φ , $\|\cdot\|_\infty$ here denotes the ℓ_∞ operator norm (maximum ℓ_1 norm of a row), $\|\cdot\|_1$ here denotes the elementwise ℓ_1 norm, and $\tilde{\gamma} \in (0, 1]$ is a constant. This condition can be seen as requiring bounded correlation between each column in the submatrix Φ_{S^c} and all columns in the submatrix Φ_S .

Condition A2. This condition is similar to requiring that $f_{\epsilon_{\ell k}}(x)$ be Lipschitz, over some x in a neighborhood of 0. We can show that the Laplace distribution, *e.g.*, satisfies this condition.

The density and distribution functions for the Laplace distribution with location zero and unit scale are given by

$$f_{\epsilon_{\ell k}}(x) = (1/2) \exp(-|x|)$$

and

$$F_{\epsilon_{\ell k}}(x) = \begin{cases} 1 - (1/2) \exp(-x) & \text{if } x \geq 0 \\ (1/2) \exp(x) & \text{if } x < 0, \end{cases}$$

respectively.

Now, suppose $0 \leq x \leq C_2$. Then we can express condition A2 as

$$|f_{\epsilon_{\ell k}}(x) - f_{\epsilon_{\ell k}}(0) - x f_{\epsilon_{\ell k}}'(0)| \leq C_1 x^2 \iff -2C_1 x^2 \leq \exp(-x) + x - 1 \leq 2C_1 x^2.$$

For the first inequality, since $1 - x \leq \exp(-x)$, it is sufficient to check that $0 \leq C_1 x^2$, which is true for $C_1 > 0$ and all x . For the second inequality, by differentiating and again using $1 - x \leq \exp(-x)$, we have that the function

$$2C_1 x^2 - \exp(-x) - x + 1 \tag{S.6}$$

is nondecreasing in $x \geq 0$; thus, it is sufficient to check that this function is nonnegative for $x = 0$, which is true.

Now, suppose $-C_2 \leq x < 0$. Then we can express condition A2 as

$$|f_{\epsilon_{\ell k}}(x) - f_{\epsilon_{\ell k}}(0) - x f_{\epsilon_{\ell k}}'(0)| \leq C_1 x^2 \iff -2C_1 x^2 \leq \exp(x) - x - 1 \leq 2C_1 x^2.$$

By symmetry with the preceding case, the first inequality here holds. The second inequality here also holds, since $\exp(x) - 2C_1x^2 - x - 1$ is continuous and increasing in $x < 0$; taking the limit as $x \uparrow 0$ gives that this function is nonpositive as required.

Condition A3. This condition is a generalization of the minimum eigenvalue condition of Wainwright (2009), *i.e.*, $c_{\min} \leq \lambda_{\min}((1/n)\Phi_S^T\Phi_S)$, for some constant $c_{\min} > 0$, and where we write Φ_S to extract the appropriate block of columns of Φ .

Condition A4. This condition allows the number of basis functions m in the expansion to grow with n , at a polynomial rate (with fractional exponent). This is roughly in line with standard nonparametric regression; *e.g.*, when estimating a continuous differentiable function via a spline expansion, one typically takes the number of basis functions m to scale as $n^{1/3}$, and the more derivatives that are assumed, the smaller the fractional exponent (Györfi et al., 2002). The condition also restricts, for any given variable, the number of variables s that contribute to its neighborhood model to be polynomial in n (with a smaller fractional exponent).

Finally, the condition assumes that the entries of the basis matrix Φ (*i.e.*, the matrix of transformed variables) to be at least of constant order, and at most of polynomial order (with small fractional exponent), with n . We note that this implicitly places a restriction on the tails of distribution governing the data $y_j^{(i)}$, $i = 1, \dots, n$, $j = 1, \dots, d$. However, the restriction is not a strong one, because it allows the maximum to grow polynomially large with n (whereas a logarithmic growth would be expected, *e.g.*, for normal data). Furthermore, it is possible to trade off the restrictions on m , s , Φ_{\max} , and d (presented in the statement of the theorem), making each of these restrictions more or less stringent, if required.

S.4 Proof of Theorem 4.3

The general strategy that we use here for support recovery is inspired by that in Fan et al. (2014), for ℓ_1 -penalized quantile regression.

Fix some $k = 1, \dots, d$ and $\ell = 1, \dots, r$. We consider the conditional distribution $y_k|y_{-k}$, whose α_ℓ -quantile is assumed to satisfy (3). Hence, to be perfectly clear, all expectations and probability statements in what follows are to be interpreted with respect to the observations $y_k^{(i)}$, $i = 1, \dots, n$ conditional on $y_j^{(i)}$, $i = 1, \dots, n$, for $j \neq k$ (and thus we can treat the feature matrix Φ as fixed throughout). In the setting assumed by the theorem, the conditional quantile model in (3) is, more explicitly,

$$Q_{y_k|y_{-k}}(\alpha_\ell) = b_{\ell k}^* + \sum_{j \neq k}^d (\theta_{\ell k j}^*)^T \phi^j(y_j),$$

for some unknown parameters $b_{\ell k}^*$ and $\theta_{\ell k j}^*$, $j = 1, \dots, d$. For simplicity, in this proof, we will drop the intercept term completely both from the model (denoted $b_{\ell k}^*$) and the optimization problem in (4) (here denoted $b_{\ell k}$) that defines the estimator in question. Including the intercept is not at all difficult, and it just requires some extra bookkeeping at various places. Recall that we define

$$S_{\ell k} = \{j \in \{1, \dots, d\} : \theta_{\ell k j}^* \neq 0\},$$

and analogously define

$$\hat{S}_{\ell k} = \{j \in \{1, \dots, d\} : \hat{\theta}_{\ell k j} \neq 0\},$$

where $\hat{\theta}_{\ell k} = (\hat{\theta}_{\ell k 1}, \dots, \hat{\theta}_{\ell k d}) \in \mathbf{R}^{dm}$ is the solution in (5).

We will show that, with probability at least $1 - \delta/(dr)$, it holds that $S_{\ell k} = \hat{S}_{\ell k}$. A union bound (over all choices $k = 1, \dots, d$ and $\ell = 1, \dots, r$) will then tell us that $E^* = \hat{E}$ with probability at least $1 - \delta$, completing the proof.

To certify that $S_{\ell k} = \hat{S}_{\ell k}$, we will show that the unique solution in (5) is given by

$$\hat{\theta}_{\ell k(S_{\ell k})} = \tilde{\theta}_{\ell k(S_{\ell k})}, \quad \hat{\theta}_{\ell k(S_{\ell k}^c)} = 0, \quad (\text{S.7})$$

where $\tilde{\theta}_{\ell k(S_{\ell k})}$ solves the “restricted” optimization problem:

$$\underset{\theta_{\ell k(S_{\ell k})}}{\text{minimize}} \quad \psi_{\alpha_\ell}(Y_k - \Phi_{S_{\ell k}} \theta_{\ell k(S_{\ell k})}) + \lambda_1 \sum_{j \in S_{\ell k}} \|\theta_{\ell k j}\|_2 + \frac{\lambda_2}{2} \|\theta_{\ell k(S_{\ell k})}\|_2^2. \quad (\text{S.8})$$

Now, to prove that $\hat{\theta}_{\ell k}$ as defined above in (S.7) indeed the solution in (5), we need to check that it satisfies the KKT conditions for (5), namely

$$\Phi_{S_{\ell k}}^T v_\ell(Y_k - \Phi_{S_{\ell k}} \tilde{\theta}_{\ell k(S_{\ell k})}) - \lambda_2 \tilde{\theta}_{\ell k(S_{\ell k})} = \lambda_1 u_{\ell k(S_{\ell k})}, \quad (\text{S.9})$$

$$\Phi_{S_{\ell k}^c}^T v_\ell(Y_k - \Phi_{S_{\ell k}} \tilde{\theta}_{\ell k(S_{\ell k})}) = \lambda_1 u_{\ell k(S_{\ell k}^c)}, \quad (\text{S.10})$$

where $v_\ell(Y_k - \Phi_{S_{\ell k}} \tilde{\theta}_{\ell k(S_{\ell k})}) \in \mathbf{R}^n$ is a subgradient of $\psi_{\alpha_\ell}(\cdot)$ at $Y_k - \Phi_{S_{\ell k}} \tilde{\theta}_{\ell k(S_{\ell k})}$, *i.e.*,

$$\left[v_\ell(Y_k - \Phi_{S_{\ell k}} \tilde{\theta}_{\ell k(S_{\ell k})}) \right]_i = \alpha_\ell - I_-(y_k^{(i)} - \Phi_{i(S_{\ell k})} \tilde{\theta}_{\ell k(S_{\ell k})}), \quad i = 1, \dots, n$$

where $I_-(\cdot)$ is the indicator function of the nonpositive real line, and where each $u_{\ell k j} \in \mathbf{R}^m$ is a subgradient of $\|\cdot\|_2$ at $\tilde{\theta}_{\ell k j}$, *i.e.*,

$$u_{\ell k j} \in \begin{cases} \{\tilde{\theta}_{\ell k j} / \|\tilde{\theta}_{\ell k j}\|_2\} & \text{if } \theta_{\ell k j} \neq 0 \\ \{x \in \mathbf{R}^m : \|x\|_2 \leq 1\} & \text{if } \theta_{\ell k j} = 0, \end{cases}$$

for $j = 1, \dots, d$. Note that, since $\tilde{\theta}_{\ell k(S_{\ell k})}$ is optimal for the restricted problem (S.8), we know that there exists a collection of subgradients $u_{\ell k(S_{\ell k})}$ to satisfy (S.9), from the KKT conditions for (S.8) itself.

It remains to satisfy (S.10), and for this, we can use $u_{\ell k j} = \Phi_j^T v_\ell(Y_k - \Phi_{S_{\ell k}} \tilde{\theta}_{\ell k(S_{\ell k})})$ as a valid choice of subgradient, for each $j \in S_{\ell k}^c$, provided that

$$\left\| \Phi_j^T v_\ell(Y_k - \Phi_{S_{\ell k}} \tilde{\theta}_{\ell k(S_{\ell k})}) \right\|_2 < \lambda_1, \quad \text{for } j \in S_{\ell k}^c. \quad (\text{S.11})$$

Define $z_j(\vartheta) = \Phi_j^T v_\ell(Y_k - \Phi_{S_{\ell k}} \vartheta)$, for $j \in S_{\ell k}^c$, and define a ball

$$B^* = \{\vartheta \in \mathbf{R}^{sm} : \|\vartheta - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq \gamma\},$$

where we write $s = |S_{\ell k}|$. To show (S.11), then, it suffices to show that

$$\underbrace{\tilde{\theta}_{\ell k(S_{\ell k})} \in B^*}_{E_1}, \quad \text{and} \quad \underbrace{\max_{j \in S_{\ell k}^c} \sup_{\vartheta \in B^*} \|z_j(\vartheta)\|_2 < \lambda_1}_{E_2}. \quad (\text{S.12})$$

In Lemma S.5.1, given in Section S.5, it is shown that the event E_1 defined above occurs with probability at least $1 - \delta/(2dr)$, with a choice of radius

$$\gamma = C \left(\frac{\lambda_1 s \sqrt{m}}{n} + \sqrt{\frac{s \log n}{n}} \right),$$

for a constant $C > 0$. Below we show that E_2 occurs with probability at least $1 - \delta/(2dr)$, as well.

For $j = 1, \dots, d$, let us expand

$$\begin{aligned} z_j(\vartheta) = & \underbrace{\Phi_j^T v_\ell(\epsilon_{\ell k})}_{\Delta_1^j} + \underbrace{\Phi_j^T \mathbf{E} \left[v_\ell(Y_k - \Phi_{S_{\ell k}} \vartheta) - v_\ell(\epsilon_{\ell k}) \right]}_{\Delta_2^j} + \\ & \underbrace{\Phi_j^T \left(v_\ell(Y_k - \Phi_{S_{\ell k}} \vartheta) - v_\ell(\epsilon_{\ell k}) - \mathbf{E} \left[v_\ell(Y_k - \Phi_{S_{\ell k}} \vartheta) - v_\ell(\epsilon_{\ell k}) \right] \right)}_{\Delta_3^j}, \quad (\text{S.13}) \end{aligned}$$

where $\epsilon_{\ell k} = (\epsilon_{\ell k 1}, \dots, \epsilon_{\ell k n}) \in \mathbf{R}^n$ is a vector of the effective error terms, which recall, is defined by $\epsilon_{\ell k} = Y_k - \Phi \theta_{\ell k}^*$. Therefore, to show that the event E_2 in (S.12) holds, we can show that for each $p = 1, 2, 3$,

$$\max_{j \in S_{\ell k}^c} \sup_{\vartheta \in B^*} \|\Delta_p^j\|_2 < \frac{\lambda_1}{3}.$$

Further, to show that E_2 holds with probability at least $1 - \delta/(2dr)$, we can show that the above holds for $p = 1, 3$ each with probability at least $1 - \delta/(4dr)$, as the statement for $p = 2$ is deterministic. We now bound the terms $\Delta_1^j, \Delta_2^j, \Delta_3^j$ one by one.

Bounding $\|\Delta_1^j\|_2$. Fix $j \in S_{\ell k}^c$, and write

$$\Phi_j^T v_\ell(\epsilon_{\ell k}) = \left(\sum_{i=1}^n \Phi_{i, (j-1)m+1} v_\ell(\epsilon_{\ell k i}), \dots, \sum_{i=1}^n \Phi_{i, jm} v_\ell(\epsilon_{\ell k i}) \right),$$

where, as a reminder that the above quantity is a vector, we have returned momentarily to the more exhaustive notation for indexing the columns of Φ , as in the right-hand side of (S.4).

Straightforward calculations reveal that, for each $i = 1, \dots, n$, and $p = 1, \dots, m$,

$$\mathbf{E} \Phi_{i, (j-1)m+p} v_\ell(\epsilon_{\ell k i}) = 0, \quad \text{and} \quad -|\Phi_{i, (j-1)m+p}| \leq \Phi_{i, (j-1)m+p} v_\ell(\epsilon_{\ell k i}) \leq |\Phi_{i, (j-1)m+p}|.$$

Hence,

$$\begin{aligned} \mathbf{Pr} \left(\|\Phi_j^T v_\ell(\epsilon_{\ell k i})\|_2 \geq \sqrt{mt} \right) & \leq \mathbf{Pr} \left(\left| \sum_{i=1}^n \Phi_{i, (j-1)m+p} v_\ell(\epsilon_{\ell k i}) \right| \geq t, \text{ some } p = 1, \dots, m \right) \\ & \leq \sum_{p=1}^m 2 \exp \left(- \frac{t^2}{2 \sum_{i=1}^n \Phi_{i, (j-1)m+p}^2} \right) \end{aligned}$$

$$\leq 2m \exp\left(-\frac{t^2}{2n}\right).$$

Above, the first inequality used the simple fact that $\|x\|_2 \leq \sqrt{m}\|x\|_\infty$ for $x \in \mathbf{R}^m$; the second used Hoeffding's bound and the union bound; and the third used our assumption that the columns of Φ have norm at most \sqrt{n} . Therefore, taking $t = \lambda_1/(3\sqrt{m})$, we see that, by the above and the union bound,

$$\mathbf{Pr}\left(\max_{j \in S_{\ell k}^c} \|\Delta_1^j\|_2 < \frac{\lambda_1}{3}\right) \geq 1 - 2dm \exp\left(-\frac{\lambda_1^2}{18mn}\right).$$

By choosing $\lambda_1 = C' \sqrt{18mn \log(8d^2mr/\delta)}$ for a constant $C' > 0$, we see that the probability in question is at least $1 - \delta/(4dr)$, as desired.

Bounding $\|\Delta_2^j\|_2$. Recall that $F_{\epsilon_{\ell k}}(\cdot)$ is used to denote the CDF of the effective error distribution, and $f_{\epsilon_{\ell k}}(\cdot)$ is used for its density. By construction, $F_{\epsilon_{\ell k}}(0) = \alpha_\ell$. Direct calculation, using the definition of $v_\ell(\cdot)$, shows that, for any $\vartheta \in B^*$, and each $i = 1, \dots, n$,

$$\mathbf{E}\left[v_\ell(\epsilon_{\ell k}) - v_\ell(Y_k - \Phi_{S_{\ell k}}\vartheta)\right] = F_{\epsilon_{\ell k}}\left(\Phi_{S_{\ell k}}(\vartheta - \theta_{\ell k(S_{\ell k})}^*)\right) - F_{\epsilon_{\ell k}}(\mathbf{0}),$$

where we apply $F_{\epsilon_{\ell k}}$ componentwise, and so

$$\Phi_j^T \mathbf{E}\left[v_\ell(\epsilon_{\ell k}) - v_\ell(Y_k - \Phi_{S_{\ell k}}\vartheta)\right] = f_{\epsilon_{\ell k}}(0)\Phi_j^T \Phi_{S_{\ell k}}(\vartheta - \theta_{\ell k(S_{\ell k})}^*) + \Delta_4^j$$

with $\Delta_4^j \in \mathbf{R}^m$ being the appropriate remainder term, *i.e.*,

$$\left[\Delta_4^j\right]_t = \sum_{i=1}^n \Phi_{it} \left[F_{\epsilon_{\ell k}}\left(\Phi_{i(S_{\ell k})}(\vartheta - \theta_{\ell k(S_{\ell k})}^*)\right) - F_{\epsilon_{\ell k}}(0) - f_{\epsilon_{\ell k}}(0)\Phi_{i(S_{\ell k})}(\vartheta - \theta_{\ell k(S_{\ell k})}^*)\right],$$

for $t = j(m-1) + 1, \dots, jm$.

Now, we have that

$$\|f_{\epsilon_{\ell k}}(0)\Phi_j^T \Phi_{S_{\ell k}}(\vartheta - \theta_{\ell k(S_{\ell k})}^*)\|_2 \leq f_{\epsilon_{\ell k}}(0)\|\Phi_j^T \Phi_{S_{\ell k}}\|_F \|\vartheta - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq \frac{\lambda_1}{6},$$

where we have used $\|\vartheta - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq \gamma$ and the groupwise irrepresentability condition in A1.

We also have the following two facts, which we will use momentarily:

$$\Phi_{\max}^3 ns\gamma^2 = o(\lambda_1) \tag{S.14}$$

$$\sqrt{s}\Phi_{\max}\gamma \rightarrow 0. \tag{S.15}$$

Note that (S.14) can be obtained as follows. Since $(1/2)(x+y)^2 \leq x^2 + y^2$ for $x, y \in \mathbf{R}$, we can plug in

$$\gamma = C \left(\frac{\lambda_1 s \sqrt{m}}{n} + \sqrt{\frac{s \log n}{n}} \right),$$

and check that both terms on the right-hand side of

$$\frac{\Phi_{\max}^3 ns}{\lambda_1} \left(\frac{\lambda_1^2 s^2 m}{n^2} + \frac{s \log n}{n} \right) = \frac{\Phi_{\max}^3 s^3 \lambda_1 m}{n} + \frac{\Phi_{\max}^3 s^2 \log n}{\lambda_1}$$

tend to zero. For the first term on the right-hand side, it is enough to show that

$$\Phi_{\max}^6 s^6 m^3 \log(d^2 mr) (\log^3 n) / n \rightarrow 0,$$

where we have plugged in $\lambda_1 = C' \sqrt{mn \log(d^2 mr / \delta) \log^3 n}$. Using the assumptions in condition A4, we get that $\log(d^2 mr) = O(\log d + \log m) = O(n^{2/21})$, and furthermore that

$$\Phi_{\max}^6 s^6 m^3 \log(d^2 mr) (\log^3 n) / n = o\left(\frac{n^{1/3} \cdot n^{2/21} \cdot n^{6/21} \cdot n^{6/21}}{\log^3 n}\right) \frac{\log^3 n}{n} \rightarrow 0,$$

as required. A similar calculation shows that the second term on the right-hand side also tends to zero, *i.e.*, $\Phi_{\max}^3 s^2 (\log n) / \lambda_1 \rightarrow 0$, which establishes (S.14). Lastly, (S.15) follows since its left-hand side is dominated by the left-hand side of (S.14).

So, we now compute

$$\begin{aligned} \|\Delta_4^j\|_2 &\leq \sqrt{m} \max_t \sum_{i=1}^n \left| \Phi_{it} \left[F_{\epsilon_{\ell k}} \left(\Phi_{i(S_{\ell k})} (\vartheta - \theta_{\ell k(S_{\ell k})}^*) \right) - \right. \right. \\ &\quad \left. \left. F_{\epsilon_{\ell k}}(0) - f_{\epsilon_{\ell k}}(0) \Phi_{i(S_{\ell k})} (\vartheta - \theta_{\ell k(S_{\ell k})}^*) \right] \right| \\ &\leq C_1 \Phi_{\max} \sqrt{m} \sum_{i=1}^n \left(\Phi_{i(S_{\ell k})} (\vartheta - \theta_{\ell k(S_{\ell k})}^*) \right)^2 \\ &\leq C_1 \Phi_{\max} \sqrt{m} \sum_{i=1}^n \|\Phi_{i(S_{\ell k})}\|_2^2 \|\vartheta - \theta_{\ell k(S_{\ell k})}^*\|_2^2 \\ &\leq C_1 \Phi_{\max}^3 \sqrt{m} n s \gamma^2 \\ &= o(\lambda_1). \end{aligned}$$

Here the first inequality follows from the fact that $\|x\|_2 \leq \sqrt{m} \|x\|_\infty$ for $x \in \mathbf{R}^m$, and the triangle inequality; the second follows from the distributional smoothness condition in A2, which is applicable since (S.15) holds; the third uses Cauchy-Schwarz; the fourth uses our column norm assumption on Φ , and $\|\vartheta - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq \gamma$; the last uses (S.14). As $\|\Delta_4^j\|_2 = o(\lambda_1)$, it will certainly be strictly less than $\lambda_1/6$ for n large enough. We have hence shown, noting that none of our above arguments have depended on the particular choice of $j = 1, \dots, d$ or $\vartheta \in B^*$,

$$\max_{j \in S_{\ell k}^c} \sup_{\vartheta \in B^*} \|\Delta_2^j\|_2 < \frac{\lambda_1}{3}.$$

Bounding $\|\Delta_3^j\|_2$. For this part, we can use the end of the proof of Lemma 2 in Fan et al. (2014), which uses classic entropy-based techniques to establish a bound very similar to that which we are seeking. By carefully looking at the conditions required for this lemma, we see that under the distributional smoothness condition in A2, condition A3, and also

$$\begin{aligned} \sqrt{n \log(dm)} &= o(\lambda_1) \\ n \Phi_{\max} \gamma^2 &= o(\lambda_1) \\ (1 + \gamma \Phi_{\max}^2 s^{3/2}) \log^2 n &= o(\lambda_1^2/n), \end{aligned}$$

all following directly from condition A4 by calculations similar to the ones we used when bounding $\|\Delta_2^j\|$, we have

$$\Pr \left(\max_{j \in S_{\ell k}^c} \sup_{\vartheta \in B^*} \|\Delta_3^j\|_2 \geq \frac{\lambda_1}{3} \right) \leq \Pr \left(\max_{j \in S_{\ell k}^c} \sup_{\vartheta \in B^*} \|\Delta_3^j\|_\infty \geq \frac{\lambda_1}{3\sqrt{m}} \right);$$

the probability on the right-hand side can be made arbitrarily small for large n , by the arguments at the end of Lemma 2 in Fan et al. (2014), and hence clearly smaller than the desired $\delta/(4dr)$ level.

Putting it together. Returning to the logic in (S.11), (S.12), (S.13), we have shown that the subgradient condition in (S.11) holds with probability at least $1 - (\delta/(2dr) + \delta/(4dr) + \delta/(4dr)) = 1 - \delta/(dr)$. Taking a union bound over $k = 1, \dots, d$ and $\ell = 1, \dots, r$, which were considered fixed at the start of our analysis, gives the result stated in the theorem. \square

S.5 Statement and proof of Lemma S.5.1

We show that with probability at least $1 - \delta/(2dr)$, it holds that $\tilde{\theta}_{\ell k(S_{\ell k})} \in B^*$, where $\tilde{\theta}_{\ell k(S_{\ell k})}$ is the solution to the restricted problem (S.8), for some fixed $k = 1, \dots, d$ and $\ell = 1, \dots, r$, and B^* is a ball defined in the proof of Theorem 4.3 in Section S.4. This fact is used a few times in the proof of Theorem 4.3.

Lemma S.5.1. *Fix some $k = 1, \dots, d$ and $\ell = 1, \dots, r$. Define the ball*

$$B^* = \{\vartheta \in \mathbf{R}^{sm} : \|\vartheta - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq \gamma\}$$

centered at the underlying coefficients $\theta_{\ell k(S_{\ell k})}^$ with radius*

$$\gamma = C \left(\frac{\lambda_1 s \sqrt{m}}{n} + \sqrt{\frac{s \log n}{n}} \right),$$

for some constant $C > 0$. Then, with probability at least $1 - \delta/(2dr)$, it holds that $\tilde{\theta}_{\ell k(S_{\ell k})} \in B^$, where $\tilde{\theta}_{\ell k(S_{\ell k})}$ is the solution to the restricted problem (S.8).*

Proof. We will follow the strategy for the proof of Theorem 1 in Fan et al. (2014) closely. We begin by considering the ball

$$B = \{\vartheta \in \mathbf{R}^{sm} : \|\vartheta - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R\}$$

with center $\theta_{\ell k(S_{\ell k})}^*$ and radius R . We also introduce some useful notational shorthand, and write the quantile loss term in the restricted problem (S.8) as

$$L_{\ell k}(\vartheta) = \psi_{\alpha_\ell}(Y_k - \Phi_{S_{\ell k}} \vartheta).$$

Below, we show that a particular function of R serves as an upper bound for the quantity $\mathbf{E}[L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*)]$, where the expectation here is taken over draws of the data, and $\tilde{\vartheta}_{\ell k(S_{\ell k})}$ is a particular point in B that we define in a moment. This in turn implies, with probability at least $1 - \delta/(2dr)$, that $\tilde{\theta}_{\ell k(S_{\ell k})} \in B^*$, as claimed.

First, we define $\tilde{\vartheta}_{\ell k(S_{\ell k})}$ more precisely: it is a point on the line segment between the solution to the restricted problem $\tilde{\theta}_{\ell k(S_{\ell k})}$ and the underlying coefficients $\theta_{\ell k(S_{\ell k})}^*$, *i.e.*,

$$\tilde{\vartheta}_{\ell k(S_{\ell k})} = \beta \tilde{\theta}_{\ell k(S_{\ell k})} + (1 - \beta) \theta_{\ell k(S_{\ell k})}^*,$$

for a particular choice

$$\beta = \frac{R}{R + \|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2},$$

which guarantees that $\tilde{\vartheta}_{\ell k(S_{\ell k})} \in B$ even if $\tilde{\theta}_{\ell k(S_{\ell k})} \notin B$, as we establish next. Observe that we always have

$$\begin{aligned} & \|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R + \|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \\ \iff & R \frac{\|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2}{R + \|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2} \leq R \\ \iff & \beta \|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R \\ \iff & \|\beta \tilde{\theta}_{\ell k(S_{\ell k})} - \beta \theta_{\ell k(S_{\ell k})}^* + \theta_{\ell k(S_{\ell k})}^* - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R \\ \iff & \|\tilde{\vartheta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R, \end{aligned}$$

as claimed. The second line here follows by rearranging and multiplying through by R ; the third by using the definition of β above; the fourth by adding and subtracting the underlying coefficients; and the fifth by using the definition of $\tilde{\vartheta}_{\ell k(S_{\ell k})}$.

Now, the beginning of the proof of Theorem 1 in Fan et al. (2014) establishes, for any $\tilde{\vartheta}_{\ell k(S_{\ell k})} \in B$, for some constant $C_5 > 0$, and using condition A3, that

$$\mathbf{E} \left[L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) \right] \geq C_5 n \|\tilde{\vartheta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2^2, \quad (\text{S.16})$$

and so, by direct calculation, since

$$\|\tilde{\vartheta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R \iff \beta \|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R \iff \|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq R/2, \quad (\text{S.17})$$

it suffices to obtain a suitable upper bound for $\mathbf{E}[L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*)]$, in order to get the result in the statement of the lemma. To this end, we introduce one more piece of shorthand, and denote the objective for the restricted problem (S.8) as $J_{\ell k}(\vartheta)$.

We proceed with the following chain of (in)equalities:

$$\begin{aligned} & \mathbf{E} \left[L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) \right] \\ &= \mathbf{E} \left[L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) \right] + J_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - J_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) + \\ & \quad J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) - J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) \quad (\text{S.18}) \\ &= \underbrace{L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) - \mathbf{E} L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) - L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) + \mathbf{E} L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})})}_{\Delta(\theta_{\ell k(S_{\ell k})}^*, \tilde{\vartheta}_{\ell k(S_{\ell k})})} + \end{aligned}$$

$$J_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) + \lambda_1 \sum_{j \in S_{\ell k}} \|\theta_{\ell k j}^*\|_2 - \lambda_1 \sum_{j \in S_{\ell k}} \|\tilde{\vartheta}_{\ell k j}\|_2 \\ - (\lambda_2/2) \|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2^2 + (\lambda_2/2) \|\theta_{\ell k(S_{\ell k})}^*\|_2^2 \quad (\text{S.19})$$

$$\leq \Delta(\theta_{\ell k(S_{\ell k})}^*, \tilde{\vartheta}_{\ell k(S_{\ell k})}) + J_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) + \lambda_1 \sum_{j \in S_{\ell k}} \|\theta_{\ell k j}^* - \tilde{\vartheta}_{\ell k j}\|_2 + o(1) \quad (\text{S.20})$$

$$\leq \Delta(\theta_{\ell k(S_{\ell k})}^*, \tilde{\vartheta}_{\ell k(S_{\ell k})}) + J_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) + \lambda_1 s R \sqrt{m} + o(1) \quad (\text{S.21})$$

$$\leq \Delta(\theta_{\ell k(S_{\ell k})}^*, \tilde{\vartheta}_{\ell k(S_{\ell k})}) + 2\lambda_1 s R \sqrt{m} \quad (\text{S.22})$$

$$\leq \sup_{\tilde{\vartheta}_{\ell k(S_{\ell k})} \in B} |\Delta(\theta_{\ell k(S_{\ell k})}^*, \tilde{\vartheta}_{\ell k(S_{\ell k})})| + 2\lambda_1 s R \sqrt{m}. \quad (\text{S.23})$$

Here, (S.18) follows by adding and subtracting like terms, and (S.19) by rearranging (S.18). In (S.20) we use the triangle inequality and the following argument to show that the terms involving λ_2 are $o(1)$. Under the assumption that $\lambda_2 = o(n^{41/42}/\theta_{\max}^*)$, combined with the restriction that $s = o(n^{1/21})$, we have $\lambda_2 = o(n/(\sqrt{s}\theta_{\max}^*))$. Therefore, under our choice of $R = 1/n$ (as specified below), we have

$$\lambda_2 \sqrt{s} \theta_{\max}^* R \rightarrow 0.$$

This in turn is used to argue that

$$-(\lambda_2/2) \|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2^2 + (\lambda_2/2) \|\theta_{\ell k(S_{\ell k})}^*\|_2^2 = (\lambda_2/2) \|\tilde{\vartheta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2^2 \\ - \lambda_2 \|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2^2 + \lambda_2 \tilde{\vartheta}_{\ell k(S_{\ell k})}^T \theta_{\ell k(S_{\ell k})}^* \\ \leq (\lambda_2/2) R^2 - \lambda_2 \|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2 (\|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2 - \|\theta_{\ell k(S_{\ell k})}^*\|_2) \\ \leq (\lambda_2/2) R^2 + \lambda_2 \|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2 R \\ \leq (\lambda_2/2) R^2 + \lambda_2 \|\theta_{\ell k(S_{\ell k})}^*\|_2 R \\ \leq (\lambda_2/2) R^2 + \lambda_2 \sqrt{s} \theta_{\max}^* R \rightarrow 0.$$

In the second to last line, we have applied $\|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2 \leq \|\theta_{\ell k(S_{\ell k})}^*\|_2$, as, outside of this case, the term in question $-(\lambda_2/2) \|\tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2^2 + (\lambda_2/2) \|\theta_{\ell k(S_{\ell k})}^*\|_2^2$ would be negative, anyway.

Continuing on, (S.21) holds because $\|\theta_{\ell k(S_{\ell k})}^* - \tilde{\vartheta}_{\ell k(S_{\ell k})}\|_2 \leq R$ implies $\|\theta_{\ell k j}^* - \tilde{\vartheta}_{\ell k j}\|_2 \leq R$. Finally, (S.22) follows because of the following argument. Since $J_{\ell k}$ is convex, we can use the definition of $\tilde{\vartheta}_{\ell k(S_{\ell k})}$ and get

$$J_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) \leq \beta J_{\ell k}(\tilde{\theta}_{\ell k(S_{\ell k})}) + (1-\beta) J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) = J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) + \beta (J_{\ell k}(\tilde{\theta}_{\ell k(S_{\ell k})}) - J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*));$$

notice that the last term here is nonpositive, since $\tilde{\theta}_{\ell k(S_{\ell k})}$ is the solution to the restricted problem (S.8), and thus we have that

$$J_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) \leq J_{\ell k}(\theta_{\ell k(S_{\ell k})}^*),$$

which lets us move from (S.21) to (S.22).

Lemma 1 in Fan et al. (2014) states, with probability at least $1 - \delta$, where $\delta = \exp(-C_6 s \log n)$ and $C_6 > 0$ is some constant, that

$$\sup_{\tilde{\vartheta}_{\ell k(S_{\ell k})} \in B} |\Delta(\theta_{\ell k(S_{\ell k})}^*, \tilde{\vartheta}_{\ell k(S_{\ell k})})| \leq 6R\sqrt{sn \log n},$$

so from (S.23), with probability at least $1 - \delta$, we see that

$$\mathbf{E} \left[L_{\ell k}(\tilde{\vartheta}_{\ell k(S_{\ell k})}) - L_{\ell k}(\theta_{\ell k(S_{\ell k})}^*) \right] \leq 6R\sqrt{sn \log n} + 2\lambda_1 s R \sqrt{m}$$

and, using (S.16), that

$$n \|\tilde{\vartheta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2^2 \leq C' \left(R\sqrt{sn \log n} + \lambda_1 s R \sqrt{m} \right),$$

for some constant $C' > 0$.

Plugging in $R = 1/n$, dividing through by n , and using the fact that the square root function is subadditive, we get, with probability at least $1 - \delta$, that

$$\begin{aligned} \|\tilde{\vartheta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 &\leq C' \left(\frac{(s \log n)^{1/4}}{n^{3/4}} + \frac{(\lambda_1 s)^{1/2} m^{1/4}}{n} \right) \\ &\leq C' \left(\sqrt{\frac{s \log n}{n}} + \frac{\lambda_1 s \sqrt{m}}{n} \right). \end{aligned}$$

Finally, we complete the proof by applying (S.17), in order to get that

$$\|\tilde{\theta}_{\ell k(S_{\ell k})} - \theta_{\ell k(S_{\ell k})}^*\|_2 \leq \gamma,$$

where we have defined

$$\gamma = C \left(\frac{\lambda_1 s \sqrt{m}}{n} + \sqrt{\frac{s \log n}{n}} \right),$$

and $C > 0$ is some constant, with probability at least $1 - \delta/(2dr)$, for large enough n . \square

S.6 Proof of Lemma 5.1

The prox operator $\mathbf{prox}_{\lambda\psi_A}(A)$ is separable in the entries of its minimizer X , so we focus on minimizing over X_{ij} the expression

$$\begin{aligned} &\max\{\alpha_j X_{ij}, (\alpha_j - 1)X_{ij}\} + (1/(2\lambda)) (X_{ij} - A_{ij})^2 \\ &= \alpha_j \max\{0, X_{ij}\} + (1 - \alpha_j) \max\{0, -X_{ij}\} + (1/(2\lambda)) (X_{ij} - A_{ij})^2. \end{aligned} \quad (\text{S.24})$$

Suppose $X_{ij} > 0$. Then differentiating (S.24) gives $X_{ij} = A_{ij} - \lambda\alpha_j$ and the sufficient condition $A_{ij} > \lambda\alpha_j$. Similarly, assuming $X_{ij} < 0$ gives $X_{ij} = A_{ij} + \lambda(1 - \alpha_j)$ when $A_{ij} < \lambda(\alpha_j - 1)$. Otherwise, we can take $X_{ij} = 0$. Putting these cases together gives the result. \square

Algorithm 1 ADMM for the MQGM

Input: observations $y^{(1)}, \dots, y^{(n)} \in \mathbf{R}^d$, feature matrix $\Phi \in \mathbf{R}^{n \times dm}$, quantile levels \mathcal{A} , constants $\lambda_1, \lambda_2 > 0$

Output: fitted coefficients $\hat{\Theta} = (\hat{\theta}_{\ell k j}, \hat{b}_{\ell k})$

for $k = 1, \dots, d$ (in parallel, if possible) **do**

initialize $\Theta_k, B_k, V, W, Z, U_V, U_W, U_Z$

repeat

 update Θ_k using (13)

 update B_k using (13)

 update V using (10)

 update W using (11)

 update Z using (12) and Lemma 5.1

 update U_V, U_W, U_Z :

$$U_V \leftarrow U_V + (\mathbf{1} B_k^T + \Phi_k \Theta - V)$$

$$U_W \leftarrow U_W + (\Theta_k - W)$$

$$U_Z \leftarrow U_Z + (Y_k \mathbf{1}^T - \mathbf{1} B_k^T - \Phi_k \Theta - Z)$$

until converged

end for

S.7 ADMM for the MQGM

A complete description of our ADMM-based algorithm for fitting the MQGM to data is given in Algorithm 1.

S.8 Additional details on Gibbs sampling

In the MQGM, there is no analytic solution for parameters like the mean, median, or quantiles of these marginal and conditional distributions, but the pseudolikelihood approximation makes for a very efficient Gibbs sampling procedure, which we highlight in this section. As it is relevant to the computational aspects of the approach, in this subsection we will make explicit the conditional random field, where y_k depends on both y_{-k} and fixed input features x .

First, note that since we are representing the distribution of $y_k | y_{-k}, x$ via its inverse CDF, to sample from this conditional distribution we can simply generate a random $\alpha \sim \text{Uniform}(0, 1)$. We then compute

$$\begin{aligned}\hat{Q}_{y_k | y_{-k}}(\alpha_\ell) &= \phi(y)^T \theta_{\ell k} + x^T \theta_{\ell k}^x \\ \hat{Q}_{y_k | y_{-k}}(\alpha_{\ell+1}) &= \phi(y)^T \theta_{(\ell+1)k} + x^T \theta_{(\ell+1)k}^x\end{aligned}$$

for some pair $\alpha_\ell \leq \alpha \leq \alpha_{\ell+1}$ and set y_k to be a linear interpolation of the two values

$$y_k \leftarrow \hat{Q}_{y_k | y_{-k}}(\alpha_\ell) + \frac{\left(\hat{Q}_{y_k | y_{-k}}(\alpha_{\ell+1}) - \hat{Q}_{y_k | y_{-k}}(\alpha_\ell) \right) (\alpha - \alpha_\ell)}{\alpha_{\ell+1} - \alpha_\ell}.$$

This highlights the desirability of having a range of non-uniformly spaced α terms that reach values close to zero and one as otherwise we may not be able to find a pair of α 's that lower and upper bound our random sample α . However, in the case that we model a sufficient quantity of α , a reasonable approximation (albeit one that will not sample from the extreme tails) is also simply to pick a random $\alpha_\ell \in \mathcal{A}$ and use just the corresponding column $\theta_{\ell k}$ to generate the random sample.

Computationally, there are a few simple but key points involved in making the sampling efficient. First, when sampling from a conditional distribution, we can precompute $x^T \Theta_k^x$ for each k , and use these terms as a constant offset. Second, we maintain a “running” feature vector $\phi(y) \in \mathbf{R}^{dm}$, *i.e.*, the concatenation of features corresponding to each coordinate $\phi(y_k)$. Each time we sample a new coordinate y_k , we generate just the new features in the $\phi(y_k)$ block, leaving the remaining features untouched. Finally, since the Θ_k terms are sparse, the inner product $\phi(y)^T \theta_{\ell k}$ will only contain a few nonzeros terms in the sum, and will be computed more efficiently if the Θ_k are stored as a sparse matrices.

S.9 Additional details on the evaluation of fitted conditional CDFs

Here, we elaborate on the evaluation of each method’s conditional CDFs that we first presented in Section 6.1. For simplicity, we describe everything that follows in terms of the conditional CDF $y_1|y_2$ only, with everything being extended in the obvious way to other conditionals. (We omit the nonparanormal skeptic from our evaluation as it is not clear how to sample from its conditionals, due to the nature of a particular transformation that it uses.)

First, we carried out the following steps in order to compute the true (empirical) conditional CDF.

1. We drew $n = 400$ samples from the ring distribution, by following the procedure described in Section 6.1; these observations are plotted across the top row of Figure S.1.
2. We then partitioned the y_2 samples into five equally-sized bins, and computed the true empirical conditional CDF of y_1 given each bin of y_2 values.

Next, we carried out the following steps in order to compute the estimated (empirical) conditional CDFs, for each method.

3. We fitted each method to the samples obtained in step (1) above.
4. Then, for each method, we drew a sample of y_1 given *each* y_2 sample, using the method’s conditional distribution; these conditionals are plotted across the second through fifth rows of Figure S.1 (for representative values of λ_1).

Operationally, we drew samples from each method’s conditionals in the following ways.

- MQGM: we used the Gibbs sampler described in Section S.8.

- MB: we drew $y_1 \sim \mathcal{N}(\hat{\theta}_1^T y_2^{(i)}, \hat{\sigma}_{1|2}^2)$, where $\hat{\theta}_1$ is the fitted lasso regression coefficient of y_1 on y_2 ; $y_2^{(i)}$ for $i = 1, \dots, n$ is the i th observation of y_2 obtained in step (1) above; and $\hat{\sigma}_{1|2}^2 = \mathbf{var}(Y_1 - Y_2 \hat{\theta}_1)$ denotes the sample variance of the underlying error term $Y_1 - Y_2 \hat{\theta}_1$ with $Y_i = (y_i^{(1)}, \dots, y_i^{(n)}) \in \mathbf{R}^n$ collecting all observations along variable i .
- SpaceJam: we drew $y_1 \sim \mathcal{N}(\hat{\theta}_1^T \phi(y_2^{(i)}), \hat{\sigma}_{1|2}^2)$, where ϕ is a suitable basis function, and $\hat{\theta}_1$ as well as $\hat{\sigma}_{1|2}^2$ are defined in ways analogous to the neighborhood selection setup.
- GLasso: we drew $y_1 \sim \mathcal{N}(\hat{\mu}_{1|2}, \hat{\sigma}_{1|2}^2)$, where

$$\begin{aligned}\hat{\mu}_{1|2} &= \hat{\mu}_1 + \hat{\Sigma}_{12} \hat{\Sigma}_{22}^{-1} (y_2^{(i)} - \hat{\mu}_2) \\ \hat{\sigma}_{1|2}^2 &= \hat{\Sigma}_{11} - \hat{\Sigma}_{12} \hat{\Sigma}_{22}^{-1} \hat{\Sigma}_{21}\end{aligned}$$

with $\hat{\mu}_i$ denoting the sample mean of Y_i , and $\hat{\Sigma}$ denoting the estimate of the covariance matrix given by GLasso (subscripts select blocks of this matrix).

5. Finally, we partitioned the y_2 samples into five equally-sized bins (just as when computing the true conditional CDF), and computed the estimated empirical conditional CDF of y_1 given each bin of y_2 values.

Having computed the true as well as estimated conditional CDFs, we measured the goodness of fit of each method's conditional CDFs to the true conditional CDFs, by computing the total variation (TV) distance, *i.e.*,

$$(1/2) \sum_{i=1}^q \left| \hat{F}_{y_1|y_2}^{\text{method}_j}(z^{(i)}|\zeta) - \hat{F}_{y_1|y_2}^{\text{true}}(z^{(i)}|\zeta) \right|,$$

as well as the (scaled) Kolmogorov-Smirnoff (KS) statistic, *i.e.*,

$$\max_{i=1, \dots, q} \left| \hat{F}_{y_1|y_2}^{\text{method}_j}(z^{(i)}|\zeta) - \hat{F}_{y_1|y_2}^{\text{true}}(z^{(i)}|\zeta) \right|.$$

Here, $\hat{F}_{y_1|y_2}^{\text{true}}(z^{(i)}|\zeta)$ is the true empirical conditional CDF of $y_1|y_2$, evaluated at $y_1 = z^{(i)}$ and given $y_2 = \zeta$, and $\hat{F}_{y_1|y_2}^{\text{method}_j}(z^{(i)}|\zeta)$ is a particular method's ("method_{*j*}" above) estimated empirical conditional CDF, evaluated at $y_1 = z^{(i)}$ and given $y_2 = \zeta$. For each method, we averaged these TV and KS values across the method's conditional CDFs. Table 1 reports the best (across a range of tuning parameters) of these averaged TV and KS values.

S.10 Additional numerical results for modeling flu epidemics

Here, we plot samples from the marginal distributions of the percentages of flu reports at regions one, five, and ten throughout the year, which reveals the heteroskedastic nature of the data (just as in Section 6.2, for region six).




Figure S.1: Conditional distributions for MQGM, MB, GLasso, and SpaceJam, fitted to samples from the ring distribution (TIGER and Laplace’s conditionals both look similar to MB’s). First row: samples from the ring distribution, where each plot highlights the samples falling into a particular shaded bin on the y_2 axis. Second through fifth rows: conditional distributions of y_1 given y_2 for each method, where each plot conditions on the appropriate y_2 bin as highlighted in the first row. The MQGM’s conditional distributions are intuitive, appearing bimodal for bin 3, and more peaked for bins 1 and 5. MB, GLasso, and SpaceJam’s densities appear (roughly) Gaussian, as expected.




Figure S.2: Samples from the fitted marginal distributions of the weekly flu incidence rates at several regions of the U.S.; samples at larger quantile levels shaded lighter, median in darker blue.




Figure S.3: Conditional independencies recovered by the MQGM on the wind farms data; each block corresponds to a wind farm, and black indicates dependence.

S.11 Sustainable energy application

We evaluate the ability of MQGM to recover the conditional independencies between several wind farms on the basis of large-scale, hourly wind power measurements; wind power is intermittent, and thus understanding the relationships between wind farms can help farm operators plan. We obtained hourly wind power measurements from July 1, 2009 through September 14, 2010 at seven wind farms ($n = 877$, see Hong et al. (2014); Wytock and Kolter (2013); Ali et al. (2016) for details). The primary variables here encode the hourly wind power at a farm over two days (*i.e.*, 48 hours), thus $d = 7 \times 48 = 336$. Exogenous variables were used to encode forecasted wind power and direction as well as other historical measurements, for a total of $q = 3417$. We set $m = 5$ and $r = 20$. Fitting the MQGM here hence requires solving $48 \times 7 = 336$ multiple quantile regression subproblems each of dimension $((336 - 1) \times 5 + 3417) \times 20 = 101,840$. Each subproblem took roughly 87 minutes, comparable to the algorithm of Wytock and Kolter (2013).

Figure S.3 presents the recovered conditional independencies; the nonzero super- and

sub-diagonal entries suggest that at any wind farm, the previous hour's wind power (naturally) influences the next hour's, while the nonzero off-diagonal entries, *e.g.*, in the (4,6) block, uncover farms that may influence one another. Wytock and Kolter (2013), whose method placed fifth in a Kaggle competition, as well as Ali et al. (2016) report similar findings (see the left panels of Figures 7 and 3 in these papers, respectively).